Covering grids with multiplicity

Bishnoi, Anurag; Boyadzhiyska, Simona; Das, Shagnik; den Bakker, Yvonne

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

eScholarship

Title
Covering grids with multiplicity

Permalink
https://escholarship.org/uc/item/7p97h8mn

Journal
Combinatorial Theory, 3(3)

ISSN
2766-1334

Authors
Bishnoi, Anurag
Boyadzhiyska, Simona
Das, Shagnik
et al.

Publication Date
2023

DOI
10.5070/C63362787

Supplemental Material
https://escholarship.org/uc/item/7p97h8mn#supplemental

Copyright Information

Peer reviewed

COVERING GRIDS WITH MULTIPLICITY

Anurag Bishnoi[1], Simona Boyadzhiyska[*2], Shagnik Das[†3], and Yvonne den Bakker[4]

[1,4]*Delft Institute of Applied Mathematics, Technische Universiteit Delft, 2628 CD Delft, Netherlands*
A.Bishnoi@tudelft.nl , ym.den.bakker@gmail.com
[2]*School of Mathematics, University of Birmingham, Edgbaston, Birmingham, B15 2TT, U.K.*
s.s.boyadzhiyska@bham.ac.uk
[3]*Department of Mathematics, National Taiwan University, Taiwan*
shagnik@ntu.edu.tw

Abstract. Given a finite grid in $\mathbb{R}^2$, how many lines are needed to cover all but one point at least k times? Problems of this nature have been studied for decades, with a general lower bound having been established by Ball and Serra. We solve this problem for various types of grids, in particular showing the tightness of the Ball–Serra bound when one side is much larger than the other. In other cases, we prove new lower bounds that improve upon Ball–Serra and provide an asymptotic answer for almost all grids. For the standard grid $\{0, \dots, n-1\} \times \{0, \dots, n-1\}$, we prove nontrivial upper and lower bounds on the number of lines needed. To prove our results, we combine linear programming duality with some combinatorial arguments.

Keywords. Grid covering, Alon–Füredi Theorem, combinatorial geometry, linear programming

Mathematics Subject Classifications. 05B40, 52C15, 05D40

1. Introduction

Given a finite grid in $\mathbb{R}^d$, a classic problem is to determine the minimum size of a k-cover: a multiset of hyperplanes that leaves one point of the grid uncovered while covering all other points at least k times. In this paper we embark on a systematic study of this problem for two-dimensional grids and obtain the following findings, which we present in full detail in Section 1.2.

1. A general lower bound of Ball and Serra is tight for $n \times m$ grids whenever $n > (k-1)(m-1)$, but is typically a constant factor too small when n is smaller. (Theorems 1.2 and 1.4)

2. The smallest k-covers of a typical $n \times n$ grid contain $\left(\frac{3}{2} + o(1)\right)k(n-1)$ lines, while all $n \times n$ grids require at least $(10 - 4\sqrt{5} + o(1))k(n-1)$ lines. (Theorem 1.5)

3. The standard grid $\{0, 1, \ldots, n-1\}^2$, where the origin is to be left uncovered, requires between $(2 - e^{-1/2} + o(1))k(n-1)$ and $(\sqrt{2} + o(1))k(n-1)$ lines to be k-covered. Moreover, the upper bound is best possible when covering with lines of slope 0, ∞, and -1. (Theorems 1.6 and 1.7)

1.1. A brief history of hyperplane covers

A celebrated result of Alon and Füredi [AF93] in combinatorial geometry states that any multiset of hyperplanes that misses exactly one point of a d-dimensional finite grid $S_1 \times \cdots \times S_d \subseteq \mathbb{F}^d$ over an arbitrary field $\mathbb{F}$ must have size at least $\sum_{i=1}^{d}(|S_i| - 1)$. This lower bound is easily seen to be tight by taking all hyperplanes of the form $x_i - a = 0$ for $1 \leqslant i \leqslant d$ and $a \in S_i \setminus \{b_i\}$, where $(b_1, \ldots, b_n)$ is the point that is uncovered. This is a significant theorem for a few different reasons; not only did the proof of Alon and Füredi play an important role in the development of the polynomial method [Alo99, Gut16, BCPS18], but this result and its generalisations have also seen several applications in a wide variety of mathematical disciplines [Kom94, BBS10, BCPS18, ABNR22, BGG+22].

One such generalisation that has been studied by several researchers is the multiplicity version of the problem, where the points of the grid should be covered multiple times. We introduce some notation to define this problem formally.

Definition 1.1. Given finite subsets $S_1, S_2, \ldots, S_d$ of a field $\mathbb{F}$, we write $\Gamma = \Gamma(S_1, S_2, \ldots, S_d)$ for the grid $S_1 \times S_2 \times \ldots \times S_d \subseteq \mathbb{F}^d$. Note that by translation we may, and will, assume $\vec{0} \in \Gamma$. For a given integer $k \geqslant 1$, we call a multiset $\mathcal{H}$ of hyperplanes in $\mathbb{F}^d$ a k-cover of Γ if every nonzero point of Γ is contained in at least k of the hyperplanes, while $\vec{0}$ is not covered at all. We denote by $\mathrm{cov}_k(\Gamma; \mathbb{F})$ the minimum cardinality of a k-cover of Γ in $\mathbb{F}^d$. In the case $\mathbb{F} = \mathbb{R}$, we shall omit the field from the notation and simply write $\mathrm{cov}_k(\Gamma)$.

In this notation, the Alon–Füredi Theorem establishes that $\mathrm{cov}_1(\Gamma, \mathbb{F}) = \sum_{i=1}^{d}(|S_i| - 1)$ for any grid Γ over any field $\mathbb{F}$. The multiplicity extension asks for the value of $\mathrm{cov}_k(\Gamma; \mathbb{F})$ for multiplicities $k \geqslant 2$, and we can start with a few trivial observations. First, if we remove any hyperplane from a k-cover, what we are left with is still a $(k-1)$-cover, and so $\mathrm{cov}_k(\Gamma; \mathbb{F}) \geqslant \mathrm{cov}_{k-1}(\Gamma; \mathbb{F}) + 1$; that is, this extremal function is strictly increasing in k. In the other direction, since the union of a k-cover and an ℓ-cover yields a $(k+\ell)$-cover, we have $\mathrm{cov}_{k+\ell}(\Gamma; \mathbb{F}) \leqslant \mathrm{cov}_k(\Gamma; \mathbb{F}) + \mathrm{cov}_\ell(\Gamma; \mathbb{F})$, and so the function is subadditive in k. Applying these recursive inequalities repeatedly until we reach the $k = 1$ case of Alon–Füredi, we have

$$\sum_{i=1}^{d}(|S_i| - 1) + k - 1 = \mathrm{cov}_1(\Gamma; \mathbb{F}) + k - 1 \leqslant \mathrm{cov}_k(\Gamma; \mathbb{F}) \leqslant k\,\mathrm{cov}_1(\Gamma; \mathbb{F}) = k\sum_{i=1}^{d}(|S_i| - 1).$$
$$(1.1)$$

The goal, then, is to narrow the considerable gap between these bounds, and there has been much previous research on some specific cases. Predating the work of Alon and Füredi [AF93], the study of affine blocking sets in finite geometry corresponds to setting $\mathbb{F} = \mathbb{F}_q$ for some prime power q and taking $\Gamma = \mathbb{F}_q^d$. For this grid, the classic paper of Jamison [Jam77] uses the polynomial method to prove $\mathrm{cov}_1(\Gamma; \mathbb{F}_q) = d(q-1)$. Bruen [Bru92] later used the polynomial method with multiplicities to provide lower bounds for the multiplicity version, showing $\mathrm{cov}_k(\Gamma, \mathbb{F}_q) \geqslant (d+k-1)(q-1)$. This is an improvement upon (1.1), but is generally not tight [Zan02, LR14]. In [BBDM23], Bishnoi et al. proved new bounds in the case of $q = 2$ by exploiting an equivalence between k-covers and linear binary codes of minimum distance k.

Recent work of Clifton and Huang [CH20] considered this problem over $\mathbb{R}$, where the grid is the hypercube $\Gamma = \{0, 1\}^d$. For fixed dimension d and growing multiplicity k, they used linear programming to determine $\mathrm{cov}_k(\Gamma)$ asymptotically. On the other hand, when the dimension d is large with respect to the multiplicity k, they applied the polynomial method to provide general lower bounds that are tight for $k = 2$ and $k = 3$. However, they conjectured that their lower bound of $d + k + 1$ is not tight for $k \geqslant 4$, and that the true value of $\mathrm{cov}_k(\Gamma)$ is $d + \binom{k}{2}$ for all fixed k and large enough d (see [CH20, Conjecture 4.1]). A subsequent paper of Sauermann and Wigderson [SW22] determined the best bound one can obtain with the polynomial method, where one seeks the minimum possible degree of a polynomial that does not vanish at the origin but has zeroes of multiplicity k at all other points in the grid. While their result improves the Clifton–Huang lower bound to $\mathrm{cov}_k(\Gamma) \geqslant d + 2k - 3$, it still falls short of the conjectured value of $\mathrm{cov}_k(\Gamma)$ in this case, suggesting a strong separation between the algebraic and geometric problems.

Apart from these special cases, Ball and Serra [BS09, Theorem 5.3] applied the polynomial method to obtain a lower bound valid for any grid $\Gamma = \Gamma(S_1, S_2, \ldots, S_d)$ over any field $\mathbb{F}$:

$$\mathrm{cov}_k(\Gamma; \mathbb{F}) \geqslant \sum_{i=1}^{d} (|S_i| - 1) + (k-1) \max_{1 \leqslant i \leqslant d} (|S_i| - 1). \tag{1.2}$$

This extends the bound of Bruen, and is a sizeable improvement on the lower bound of (1.1), but remains far removed from the upper bound. Indeed, in the symmetric case when $|S_i| = n$ for all $i \in [d]$, we have $(d + k - 1)(n - 1) \leqslant \mathrm{cov}_k(\Gamma; \mathbb{F}) \leqslant kd(n-1)$. It is thus of great interest to determine whether the Ball–Serra bound can be tight, and to obtain better bounds when it is not.

1.2. Our results

In this paper we work in the setting of $\mathbb{R}^2$, determining $\mathrm{cov}_k(\Gamma)$ for a wide array of two-dimensional real grids. Our first set of results concerns how the dimensions of the grid affect the tightness of the Ball–Serra bound, which we show to be sharp when the grid is much wider than it is tall.

Theorem 1.2. *Let $S_1, S_2 \subseteq \mathbb{R}$ satisfy $|S_1| = n$, $|S_2| = m$, and $0 \in S_1 \cap S_2$, and set $\Gamma = \Gamma(S_1, S_2) \subseteq \mathbb{R}^2$. If, for a positive integer k, we have $n \geqslant (k-1)(m-1)+1$, then*

$$\mathrm{cov}_k(\Gamma) = k(n-1) + (m-1).$$

Our next result shows that the lower bound on n from Theorem 1.2 cannot be improved in general. In fact, when $n \leqslant (k-1)(m-1)$, the Ball–Serra bound can be improved for almost all $n \times m$ grids. We make the "almost all" precise with the following definition.

Definition 1.3. Let $S_1, S_2 \subseteq \mathbb{R}$ with $0 \in S_1 \cap S_2$, let $\Gamma = \Gamma(S_1, S_2)$ and let $\Delta \geqslant 0$ be an integer. We call a point in Γ an *axial point* if any of its coordinates is equal to 0, and a *nonaxial point* otherwise. We call Γ Δ-*generic* if any line containing two axial points contains at most Δ nonaxial points. In the case $\Delta = 0$, when such lines avoid all nonaxial points, we simply call Γ *generic*.

To see that grids are typically generic, suppose the nonzero points of S_1 and S_2 are sampled uniformly and independently from $[-1, 1]$. If a line contains two axial points and a nonaxial point, then the axial points must come from different axes, and so we may assume the points in question are $(0, b_1)$, $(a_1, 0)$ and (a_2, b_2) for some nonzero $a_1, a_2 \in S_1$ and $b_1, b_2 \in S_2$. For these points to be collinear, we require $a_1(b_1 - b_2) = a_2 b_1$, and the probability of such an equation holding for any choice of a_i, b_j, of which there are only finitely many, is 0. We are now ready to state our next result, which concerns $\mathrm{cov}_k(\Gamma)$ in the case where Γ is a generic grid.

Theorem 1.4. *Let $\Gamma = \Gamma(S_1, S_2) \subseteq \mathbb{R}^2$ be a generic grid with $|S_1| = n$, $|S_2| = m$ and $0 \in S_1 \cap S_2$. Then*

$$\mathrm{cov}_k(\Gamma) \geqslant k(n-1) + \frac{k}{n+m-2}(m-1)^2.$$

Furthermore, if Γ is an arbitrary $n \times m$ grid, and if k is divisible by $\frac{n+m-2}{\gcd(n-1, m-1)}$, then we have

$$\mathrm{cov}_k(\Gamma) \leqslant k(n-1) + \frac{k}{n+m-2}(m-1)^2. \tag{1.3}$$

In particular, if Γ is generic, we have equality above.

Note that $\frac{k}{n+m-2}(m-1)^2$ is strictly larger than $m-1$ precisely when $n \leqslant (k-1)(m-1)$, and hence these theorems show the Ball–Serra bound is tight for generic grids if and only if $n \geqslant (k-1)(m-1)+1$.

Theorem 1.4 shows that the Ball–Serra bound gets worse as the dimensions of the grid grow closer. For the majority of our paper, therefore, we focus on square grids Γ, where $n = m$. In this case, the Ball–Serra bound gives $\mathrm{cov}_k(\Gamma) \geqslant (k+1)(n-1)$. Our next theorem gives a stark improvement for general square grids.

Theorem 1.5. *Let* $\Gamma = \Gamma(S_1, S_2) \subset \mathbb{R}^2$ *be a grid with* $|S_1| = |S_2| = n$ *and* $0 \in S_1 \cap S_2$. *Then, for any integer* $k \geqslant 2$, *we have:*

(a) $\operatorname{cov}_k(\Gamma) \leqslant \left\lceil \frac{3}{2}k \right\rceil (n - 1)$.

(b) $\operatorname{cov}_k(\Gamma) \geqslant \left(10 - 4\sqrt{5} + O(n^{-1})\right)k(n - 1)$.

(c) *if* Γ *is* Δ*-generic for some* $\Delta \geqslant 0$, *then* $\operatorname{cov}_k(\Gamma) \geqslant \left(2 - \frac{n-1}{2(n-1)-\Delta}\right)k(n-1)$. *In particular, we have* $\operatorname{cov}_k(\Gamma) \geqslant \frac{3}{2}k(n-1)$ *when* Γ *is generic.*

Note that part (a) improves the trivial upper bound from (1.1), while, since $10 - 4\sqrt{5} \approx 1.0557$, part (b) gives a constant factor improvement over the Ball–Serra bound in (1.2), showing that it is never tight for large square grids. Moreover, as previously discussed, almost all grids are generic, and part (c) gives a lower bound that matches the upper bound of part (a) exactly when k is even and asymptotically when k is odd and large. In fact, the bound is asymptotically tight for all $o(n)$-generic grids.

However, the most natural grid to consider is the *standard* $n \times n$ *grid*, given by $S_1 = S_2 = \{0, 1, 2, \ldots, n - 1\}$, and we denote this by $\Gamma_n = \Gamma(S_1, S_2)$. By considering the diagonal $x + y = n - 1$, we see that Γ_n is not Δ-generic for any $\Delta < n - 2$, and so Theorem 1.5(c) is worse than the Ball–Serra bound when $n > k$. By tailoring our methods to this specific grid, we obtain the following improvements on the general bounds.

Theorem 1.6. *Let* $n, k \geqslant 2$ *be integers, let* $S = \{0, 1, 2, \ldots, n - 1\}$ *and let* $\Gamma_n = \Gamma(S, S)$ *be the standard* $n \times n$ *grid. Then we have*

$$\left(2 - e^{-1/2} + O(n^{-1})\right)k(n - 1) \leqslant \operatorname{cov}_k(\Gamma_n) \leqslant \left(\sqrt{2} + O(n^{-1} + k^{-1})\right)k(n - 1). \quad (1.4)$$

Note that $2 - e^{-1/2} \approx 1.3935$, while $\sqrt{2} \approx 1.4142$, so there is still a gap between the best lower and upper bounds we obtain for standard grids. To obtain sharper bounds, it can help to restrict the class of k-covers we consider. As we will see in the proof of the upper bound in Theorem 1.6, our construction uses only three types of lines: horizontal, vertical, and lines of slope -1. A subsequent computer search verified that for small values of n and k, we can always find an optimal k-cover using only these three kinds of lines. In our final result, we provide a matching lower bound for these restricted k-covers, suggesting the upper bound of (1.4) may be correct in the unrestricted case as well.

Theorem 1.7. *The smallest* k*-cover of the standard grid* Γ_n *that only contains lines of slope* 0, ∞, *or* -1 *has size at least* $\left(\sqrt{2} + O(n^{-1})\right)k(n - 1)$.

When proving these results, we shall establish the upper bounds by means of explicit constructions of k-covers. The lower bounds, meanwhile, will follow by applying duality to a linear programming relaxation of this problem, and we shall set up this framework in Section 2. Having described our methodology, we will prove Theorems 1.2 and 1.4 in Section 3. We then shift our focus to square grids, proving Theorems 1.5, 1.6, and 1.7 in Section 4. Finally, we provide some concluding remarks and open problems in Section 5.

2. The linear programming framework

In this section we introduce the linear programming method that will be used to prove our lower bounds. The use of linear programming in extremal combinatorics is well-established and has led to many results (see, for example, [GMCS22]), including the Clifton–Huang [CH20] lower bound on $\mathrm{cov}_k(\Gamma)$ for $\Gamma = \{0,1\}^d$ in the case when d is fixed and k tends to infinity. The standard template is as follows: assume without loss of generality that we are working with a minimization problem. First, we interpret our extremal problem as an instance of an integer programming problem. Then, since integer programming is intractable, we consider the linear programming relaxation, where we allow fractional solutions. Since an integer solution is in particular a fractional solution, the value of the linear program gives a lower bound to the original extremal problem. Crucially, we can then apply duality, and the task of finding lower bounds translates to finding feasible solutions to the dual linear program. We refer the reader to [MG07] for further background on linear programming.

To start with this plan of action, we need to recast the covering problem as an integer linear program. Given a grid $\Gamma = \Gamma(S_1, S_2)$, for some finite sets $S_1, S_2 \subseteq \mathbb{R}$ containing 0, we wish to determine $\mathrm{cov}_k(\Gamma)$, the minimum number of lines in a k-cover of Γ. We shall assume $|S_1|, |S_2| \geqslant 2$, as the problem is otherwise trivial.

For every possible line ℓ not containing the origin, we introduce a variable $z(\ell)$ indicating its multiplicity in the k-cover. We thus wish to minimise $\sum_\ell z(\ell)$, the size of the cover. In order to ensure that the solution returned by the integer program is a k-cover, we need to require that each nonzero point of Γ be covered at least k times while the origin be omitted altogether. The latter constraint is easily seen to be satisfied. Hence, for each $(x,y) \in \Gamma \setminus \{(0,0)\}$, we require that the sum of $z(\ell)$ over all origin-avoiding lines ℓ containing (x,y) be at least k.

The one catch is that there are infinitely many lines in $\mathbb{R}^2$. To obtain a finite program, we observe that we may restrict our attention to lines that contain at least two nonzero points in Γ. Indeed, in any minimal k-cover of Γ, every line must contain at least one nonzero point, and if a line contains only one point (x,y), then we can replace it by a different origin-avoiding line passing through (x,y) and at least one other point of Γ. Thus, we need only consider lines from the set $\mathcal{L} = \mathcal{L}(\Gamma)$ of origin-avoiding lines containing at least two points of Γ. Since there are fewer than $|\Gamma|^2$ pairs of nonzero points in Γ, each of which determines a unique line, it follows that $\mathcal{L}$ is finite.

In summary, $\mathrm{cov}_k(\Gamma)$ is the solution to the following integer linear program $\mathcal{I} = \mathcal{I}(\Gamma, k)$.

$$\text{minimize} \quad \sum_{\ell \in \mathcal{L}} z(\ell)$$

$$\text{subject to}$$

$$\sum_{\substack{\ell \in \mathcal{L}: \\ (x,y) \in \ell}} z(\ell) \geqslant k \quad \text{for all } (x,y) \in \Gamma \setminus \{(0,0)\}$$

$$z(\ell) \in \mathbb{Z}_{\geqslant 0} \quad \text{for all } \ell \in \mathcal{L}$$

The final constraint, that the variables $z(\ell)$ be integral, renders solving the program computationally infeasible. Instead, to obtain a polynomial-time solvable problem, we can relax the

variables to be real-valued. Now that we are no longer constrained to the integers, we can also divide through by k, removing the dependence of the program on this parameter. We therefore obtain the linear program $\mathcal{P} = \mathcal{P}(\Gamma)$ with the normalised variables $u(\ell)$ for $\ell \in \mathcal{L}$.

$$\text{minimize} \quad \sum_{\ell \in \mathcal{L}} u(\ell) \tag{2.1}$$

$$\text{subject to}$$

$$\sum_{\substack{\ell \in \mathcal{L}: \\ (x,y) \in \ell}} u(\ell) \geqslant 1 \quad \text{for all } (x,y) \in \Gamma \setminus \{(0,0)\}$$

$$u(\ell) \geqslant 0 \quad \text{for all } \ell \in \mathcal{L}$$

Let us denote by $\Phi(\Gamma)$ the solution to $\mathcal{P}(\Gamma)$. The following result shows that $\Phi(\Gamma)$ describes the asymptotic behaviour of $\operatorname{cov}_k(\Gamma)$ when k is large with respect to the dimensions of the grid.

Proposition 2.1. *For any grid Γ and integer $k \geqslant 1$ we have*

$$k\Phi(\Gamma) \leqslant \operatorname{cov}_k(\Gamma) \leqslant k\Phi(\Gamma) + |\mathcal{L}|.$$

Proof. As previously established, $\operatorname{cov}_k(\Gamma)$ is the value of the integer linear program $\mathcal{I}(\Gamma, k)$. If we let $(z(\ell) : \ell \in \mathcal{L})$ be a solution to the program, then setting $u(\ell) = \frac{1}{k}z(\ell)$ yields a feasible solution to the linear relaxation $\mathcal{P}(\Gamma)$, with value $\sum_\ell u(\ell) = \frac{1}{k}\sum_\ell z(\ell) = \frac{1}{k}\operatorname{cov}_k(\Gamma)$. As $\Phi(\Gamma)$ is the minimum possible value of a feasible solution to $\mathcal{P}(\Gamma)$, the first inequality follows.

For the upper bound, let $(u^*(\ell) : \ell \in \mathcal{L})$ be an optimal solution to the linear program $\mathcal{P}(\Gamma)$, with value $\Phi(\Gamma)$. If we then set $z(\ell) = \lceil ku^*(\ell) \rceil$ for all $\ell \in \mathcal{L}$, we obtain a feasible solution to $\mathcal{I}(\Gamma, k)$. Thus,

$$\operatorname{cov}_k(\Gamma) \leqslant \sum_\ell z(\ell) = \sum_\ell \lceil ku^*(\ell) \rceil \leqslant \sum_\ell (ku^*(\ell) + 1) = k\Phi(\Gamma) + |\mathcal{L}|. \qquad \square$$

We remark that in practice one often obtains better error bounds; for instance, the $|\mathcal{L}|$ term can be replaced by the size of the support of the optimal solution u^*. Furthermore, for an infinite sequence of multiplicities k, we can do away with the error term altogether. Indeed, since all the coefficients of $\mathcal{P}(\Gamma)$ are integral, there is a rational optimal solution u^* (the one returned by the Simplex Algorithm, for example). If k is divisible by the common denominator of the fractions $u^*(\ell)$, then we can set $z(\ell) = ku^*(\ell)$ without needing any rounding, thereby obtaining a solution to $\mathcal{I}(\Gamma, k)$ of value precisely $k\Phi(\Gamma)$.

Therefore, asymptotically as k tends to infinity, the problem reduces to determining $\Phi(\Gamma)$. We can provide upper bounds by finding feasible solutions to $\mathcal{P}(\Gamma)$ or, better yet, constructing k-covers of Γ. To obtain lower bounds, we appeal to the theory of duality. The dual of $\mathcal{P}(\Gamma)$, which we denote by $\mathcal{D}(\Gamma)$, is the following linear program, where we have a variable $w(x,y)$ for

each point $(x, y) \in \Gamma \setminus \{(0,0)\}$, which we call the weight of the point.

$$\text{maximize} \sum_{(x,y) \in \Gamma \setminus \{(0,0)\}} w(x,y) \qquad (2.2)$$

subject to

$$\sum_{\substack{(x,y) \in \Gamma \setminus \{(0,0)\}: \\ (x,y) \in \ell}} w(x,y) \leqslant 1 \quad \text{for all } \ell \in \mathcal{L}$$

$$w(x,y) \geqslant 0 \quad \text{for all } (x,y) \in \Gamma \setminus \{(0,0)\}$$

For convenience, given a set $S \subseteq \Gamma \setminus \{(0,0)\}$, we write $w(S) = \sum_{(x,y) \in S} w(x,y)$ for the weight of S. The dual program thus asks for the maximum possible weight of the grid, provided every line in $\mathcal{L}$ has weight at most 1. By the duality theorem for linear programming (see [MG07, Section 6.1]), the programs $\mathcal{P}(\Gamma)$ and $\mathcal{D}(\Gamma)$ have the same optimal objective value $\Phi(\Gamma)$. We shall thus prove our lower bounds on $\mathrm{cov}_k(\Gamma)$ by finding suitably large feasible weights on the grid Γ.

3. Wide rectangular grids

In this section we will prove Theorems 1.2 and 1.4, establishing precisely when the Ball–Serra bound is tight for all grids of given dimensions. Our first result establishes the value of $\mathrm{cov}_k(\Gamma)$ for all $n \times m$ grids Γ whenever $n \geqslant (k-1)(m-1) + 1$.

Proof of Theorem 1.2. The Ball–Serra bound (1.2) provides the requisite lower bound, as substituting $|S_1| = n$ and $|S_2| = m$ gives $\mathrm{cov}_k(\Gamma) \geqslant k(n-1) + m - 1$. To prove a matching upper bound, we provide an explicit construction of a k-cover containing this many lines.

Write $S_2 = \{0, t_1, \ldots, t_{m-1}\}$, and let $P_1 \cup \cdots \cup P_{m-1}$ be an arbitrary partition of $S_1 \setminus \{0\}$ such that $|P_i| \geqslant k - 1$ for all $i \in [m-1]$; such a partition exists since $n - 1 \geqslant (k-1)(m-1)$.

Now, consider the following collection of lines:

(i) the line $y = t_i$ for all $i \in [m-1]$;

(ii) $k - 1$ copies of the line $x = s$ for all $s \in S_1 \setminus \{0\}$;

(iii) the line connecting $(0, t_i)$ and $(s, 0)$ for every $i \in [m-1]$ and $s \in P_i$.

In total, this collection contains $m - 1 + (k-1)(n-1) + n - 1 = k(n-1) + m - 1$ lines. It remains to verify that these lines form a valid k-cover of Γ. Note first that no line in this collection passes through the origin $(0,0)$.

Any nonaxial point of Γ is covered k times by the lines in (i) and (ii), leaving us to check the axial points. A point of the form $(s, 0)$, where $s \in S_1 \setminus \{0\}$, is covered $k - 1$ times by the lines in (ii) and once by the lines in (iii). Finally, a point $(0, s)$ for $s \in S_2 \setminus \{0\}$ is covered once by the lines in (i) and at least $k - 1$ times by the lines in (iii) since each P_i has size at least $k - 1$. Hence, every nonzero point of Γ is covered at least k times, as required. □

We remark that the above construction can be generalised to higher dimensions to show that, for any $n_1 \times \cdots \times n_d$ grid $\Gamma(S_1, \ldots, S_d)$ containing the origin, if $n_1 \geqslant n_2 \geqslant \ldots \geqslant n_d$ and $n_1 \geqslant \mathrm{cov}_{k-1}(\Gamma(S_2, \ldots, S_d)) + 1$, then the Ball–Serra bound is tight for $\mathrm{cov}_k(\Gamma(S_1, \ldots, S_d))$. Indeed, write $S_1 = \{0, s_1, \ldots, s_{n_1-1}\}$ and let $\mathcal{H} = \{H_1, \ldots, H_{n_1-1}\}$ be any collection of $n_1 - 1$ hyperplanes in $\mathbb{R}^{d-1}$ containing a $(k - 1)$-cover of $\Gamma(S_2, \ldots, S_d)$. We then form a k-cover of $\Gamma(S_1, S_2, \ldots, S_d)$ consisting of the following hyperplanes in $\mathbb{R}^d$:

(i) one copy of the hyperplane $x_i = t$ for all $i \in \{2, \ldots, d\}$ and $t \in S_i \setminus \{0\}$;

(ii) $k - 1$ copies of the hyperplane $x_1 = s$ for all $s \in S_1 \setminus \{0\}$;

(iii) the hyperplane spanned by $\{0\} \times H_i$ and $(s_i, 0, \ldots, 0)$ for all $i \in [n_1 - 1]$.

It is not difficult to check that this is indeed a k-cover of $\Gamma(S_1, \ldots, S_d)$, and it consists of $\sum_{i=1}^{d} n_i + (k-1)n_1$ hyperplanes, which matches the Ball–Serra lower bound (1.2). However, it is not clear how good the lower bound on n_1 is; that is, how large n_1 needs to be with respect to the other dimensions n_i in order to ensure that the Ball–Serra bound is tight.

In our next result, we show that the bound on n_1 in Theorem 1.2 is best possible, since for generic grids that are slightly less wide, the Ball–Serra bound is no longer tight. In fact, we a give a general lower bound for $\mathrm{cov}_k(\Gamma)$ when Γ is a generic $n \times m$ grid, and prove that this bound is tight for infinitely many choices of k. While we will not pursue this question further for higher dimensions in this paper, we remark that it was shown in [dB21] that for the grid $\{0, \ldots, n_1 - 1\} \times \{0, \ldots, n_2 - 1\} \times \{0, \ldots, n_3 - 1\}$ with $n_1 \geqslant n_2 \geqslant n_3$, the Ball–Serra bound is already tight when $n_1 \geqslant (k - 1)(n_2 - 1) + 1$, which is an improvement on the bound given by the above construction.

Proof of Theorem 1.4. We wish to show that if $\Gamma = \Gamma(S_1, S_2)$ is a generic grid, where $S_1, S_2 \subseteq \mathbb{R}$ satisfy $0 \in S_1 \cap S_2$ and $|S_1| = n \geqslant m = |S_2|$, then we have $\mathrm{cov}_k(\Gamma) \geqslant k(n-1) + \frac{k}{n+m-2}(m-1)^2$. Appealing to the linear programming framework developed in Section 2, it suffices to show $\Phi(\Gamma) \geqslant (n - 1) + \frac{(m-1)^2}{n+m-2}$, which can be done by defining a weighting on the nonzero points of Γ with this total weight in which every line in $\mathcal{L}$ has weight at most 1.

To that end, define the weighting $w : \Gamma \setminus \{(0, 0)\} \to \mathbb{R}$ by

$$w((x, y)) = \begin{cases} \frac{n-1}{n+m-2} & \text{if } y = 0; \\ \frac{m-1}{n+m-2} & \text{if } x = 0; \\ \frac{1}{n+m-2} & \text{otherwise.} \end{cases}$$

We start by computing the total weight of the grid:

$$w(\Gamma \setminus \{(0,0)\}) = (n-1)\frac{n-1}{n+m-2} + (m-1)\frac{m-1}{n+m-2} + (n-1)(m-1)\frac{1}{n+m-2}$$

$$= (n-1)\frac{n-1+m-1}{n+m-2} + \frac{(m-1)^2}{n+m-2}$$

$$= n-1 + \frac{(m-1)^2}{n+m-2},$$

and so, provided this weighting is feasible, it gives the desired lower bound.

To establish its feasibility, let us consider a line $\ell \in \mathcal{L}(\Gamma)$. First suppose ℓ contains two axial points, say $(x, 0)$ and $(0, y)$. Since Γ is generic, ℓ cannot contain any other points, and hence $w(\ell) = w((x, 0)) + w((0, y)) = \frac{n-1}{n+m-2} + \frac{m-1}{n+m-2} = 1$. Next, suppose ℓ is a horizontal line of the form $y = s$ for some $s \in S_2 \setminus \{0\}$. The line ℓ then contains one point on the y-axis and $n - 1$ nonaxial points, and thus $w(\ell) = \frac{m-1}{n+m-1} + (n-1)\frac{1}{n+m-2} = 1$. Finally, any other line ℓ can contain at most one axial point and at most $m - 1$ nonaxial points (x, y), one for each choice of $y \in S_2 \setminus \{0\}$. For such lines, we therefore have $w(\ell) \leqslant \frac{n-1}{n+m-2} + (m-1)\frac{1}{n+m-2} = 1$.

Hence, the weighting w is indeed feasible for the dual linear program $\mathcal{D}(\Gamma)$, and has total weight $n - 1 + \frac{(m-1)^2}{n+m-2}$, which proves $\mathrm{cov}_k(\Gamma) \geqslant k(n-1) + \frac{k}{n+m-2}(m-1)^2$.

Now, given an arbitrary $n \times m$ grid $\Gamma = \Gamma(S_1, S_2)$, we provide a construction of a k-cover of Γ that matches the bound proven above for an infinite sequence of multiplicities k, thereby determining $\mathrm{cov}_k(\Gamma)$ for generic grids Γ and such multiplicities k.

Defining $a = \frac{n-1}{\gcd(n-1,m-1)}$ and $b = \frac{m-1}{\gcd(n-1,m-1)}$, we are given that k is divisible by $a + b$. We further define $d_1 = \frac{bk}{a+b}$ and $d_2 = \frac{ak}{a+b}$, noting that our divisibility assumption ensures these are integers and that $d_1 + d_2 = k$. Let B be an arbitrary biregular bipartite multigraph with parts $S_1 \setminus \{0\}$ and $S_2 \setminus \{0\}$ with degrees d_1 in the first part and d_2 in the second. Note that such a multigraph exists, since $d_1(n - 1) = d_2(m - 1)$, and we can assign d_1 half-edges to each $s_1 \in S_1 \setminus \{0\}$ and d_2 half-edges to each $s_2 \in S_2 \setminus \{0\}$, and then take an arbitrary matching between the two sets of half-edges.

Next, consider the following collection of lines:

(i) d_2 copies of the line $x = s_1$ for each $s_1 \in S_1 \setminus \{0\}$;

(ii) d_1 copies of the line $y = s_2$ for each $s_2 \in S_2 \setminus \{0\}$;

(iii) for each $\{s_1, s_2\} \in E(B)$, a copy of the line connecting $(s_1, 0)$ to $(0, s_2)$.

To see that these lines form a k-cover of Γ, observe that every nonaxial point is covered by d_2 vertical lines from (i) and d_1 horizontal lines from (ii), and is thus covered $d_1 + d_2 = k$ times in total. For the axial points, a point of the form $(s_1, 0)$ for $s_1 \in S_1 \setminus \{0\}$ is covered d_2 times by the lines in (i), while the biregularity of the multigraph B ensures it is covered d_1 times by the lines in (iii). Similarly, each point of the form $(0, s_2)$ for $s_2 \in S_2 \setminus \{0\}$ is covered d_1 times by the lines in (ii) and d_2 times by those in (iii). Thus, the axial points are also each covered k times. Finally, none of the lines in our collection passes through the origin.

We thus obtain our upper bound by calculating the size of this cover, which yields

$$
\begin{aligned}
\mathrm{cov}_k(\Gamma) &\leqslant d_2(n - 1) + d_1(m - 1) + d_1(n - 1) \\
&= (d_1 + d_2)(n - 1) + d_1(m - 1) \\
&= k(n - 1) + \frac{bk}{a + b}(m - 1) \\
&= k(n - 1) + \frac{k}{n + m - 2}(m - 1)^2,
\end{aligned}
$$

as required. □

4. Square grids

In the previous section, we saw that the Ball–Serra bound is tight when the grid is much wider than it is tall, and proved a lower bound for generic grids that becomes much larger than the Ball–Serra bound as the dimensions grow closer in size. Therefore, for the rest of this paper, we focus on $n \times n$ grids. In Section 4.1 we prove Theorem 1.5, which provides general lower and upper bounds on $\mathrm{cov}_k(\Gamma)$ for arbitrary $n \times n$ grids Γ. In Section 4.2 we focus on the standard grid $\Gamma_n = \Gamma(\{0, 1, \ldots, n-1\}, \{0, 1, \ldots, n-1\})$, proving Theorems 1.6 and 1.7.

4.1. General results

We start by restating our general bounds for square grids.

Theorem 1.5. *Let $\Gamma = \Gamma(S_1, S_2) \subset \mathbb{R}^2$ be a grid with $|S_1| = |S_2| = n$ and $0 \in S_1 \cap S_2$. Then, for any integer $k \geqslant 2$, we have:*

(a) $\mathrm{cov}_k(\Gamma) \leqslant \left\lceil \frac{3}{2}k \right\rceil (n-1)$.

(b) $\mathrm{cov}_k(\Gamma) \geqslant \left(10 - 4\sqrt{5} + O(n^{-1})\right) k(n-1)$.

(c) *if Γ is Δ-generic for some $\Delta \geqslant 0$, then $\mathrm{cov}_k(\Gamma) \geqslant \left(2 - \frac{n-1}{2(n-1)-\Delta}\right) k(n-1)$. In particular, we have $\mathrm{cov}_k(\Gamma) \geqslant \frac{3}{2}k(n-1)$ when Γ is generic.*

We will prove the upper bound of (a) via an explicit construction of a k-cover, and use the linear programming framework of Section 2 to establish the lower bounds of (b) and (c).

Proof. (a) Note that when k is even, the upper bound $\mathrm{cov}_k(\Gamma) \leqslant \frac{3}{2}k(n-1)$ follows from the upper bound in Theorem 1.4 when $m = n$. We will obtain the upper bound for odd k with some appropriate rounding, and present a unified construction below.

Let $\{x_1, x_2, \ldots, x_{n-1}\}$ be the nonzero elements of S_1, and let $\{y_1, y_2, \ldots, y_{n-1}\}$ be the nonzero elements of S_2. We form a k-cover of Γ consisting of the following lines:

(i) $\left\lceil \frac{1}{2}k \right\rceil$ copies of the lines $x = x_i$ and $y = y_i$, for each $i \in [n-1]$;

(ii) $\left\lfloor \frac{1}{2}k \right\rfloor$ copies of line connecting $(x_i, 0)$ to $(0, y_i)$, for each $i \in [n-1]$.

There are $2\left\lceil \frac{1}{2}k \right\rceil(n-1)$ lines in (i) and $\left\lfloor \frac{1}{2}k \right\rfloor$ lines in (ii), and hence we have a total of $\left\lceil \frac{3}{2}k \right\rceil(n-1)$ lines, and it is evident that none of these pass through the origin. To see that they form a k-cover, observe first that each nonaxial point is covered by $\left\lceil \frac{1}{2}k \right\rceil$ horizontal lines and $\left\lceil \frac{1}{2}k \right\rceil$ vertical lines, and is thus covered at least k times in total. Meanwhile, each axial point is covered $\left\lceil \frac{1}{2}k \right\rceil$ times by the lines in (i) and $\left\lfloor \frac{1}{2}k \right\rfloor$ times by the lines in (ii), and so is incident to exactly k lines. These lines therefore indeed form a k-cover, showing $\mathrm{cov}_k(\Gamma) \leqslant \left\lceil \frac{3}{2}k \right\rceil(n-1)$.

(b) To prove lower bounds, we appeal to the dual linear program $\mathcal{D}(\Gamma)$. Our goal is to define a weighting w of the nonzero points of Γ of large total weight in which no origin-avoiding line has weight more than 1. The key observation is that these lines can contain at most two axial points, and so we can hope to get away with assigning large weights to the axial points.

We shall first try a simple weighting w', where all nonaxial points obtain a weight of α, and all axial points a weight of β, for α and β to be chosen later. Unfortunately, this initial attempt does not work. Indeed, we have $w'(\Gamma) = (n-1)^2\alpha + 2(n-1)\beta = ((n-1)\alpha + 2\beta)(n-1)$. Since there could be lines containing $n-2$ nonaxial points and 2 axial points, we must have $(n-2)\alpha + 2\beta \leqslant 1$. Hence, $w'(\Gamma) \leqslant (1+\alpha)(n-1)$. As lines parallel to the axes contain $n-1$ nonaxial points and an axial point, we must also have $(n-1)\alpha + \beta \leqslant 1$, which in particular implies $\alpha \leqslant \frac{1}{n-1}$. Thus, $w'(\Gamma) \leqslant n$, and so the best lower bound we can hope for from such a weighting is $\mathrm{cov}_k(\Gamma) \geqslant kn$, which is worse than the Ball–Serra bound (1.2) for $n \geqslant k+2$.

To salvage this idea, we will instead only assign weight to some of the axial points, with the aim of ensuring that any origin-avoiding line containing two positively-weighted axial points cannot contain too many nonaxial points. For this, we use the following claim, bounding the number of points contained in certain lines.

Claim 4.1. *Suppose we enumerate the members of S_1 as $x_1 < x_2 < \ldots < x_n$ and those of S_2 as $y_1 < y_2 < \ldots < y_n$, and suppose $i_0 \in [n]$ is such that $x_{i_0} = 0$. Let $j \in [n]$ and ℓ be a line passing through $(0, y_j)$. If ℓ has positive slope, then ℓ contains at most $n - |j - i_0|$ points of Γ. If ℓ has negative slope, then ℓ contains at most $n - |(n - j + 1) - i_0|$ points of Γ.*

Proof of Claim 4.1. Suppose first that ℓ is a line of positive slope passing through $(0, y_j)$. We define $S_1^- = \{x \in S_1 : x \leqslant 0\}$ and $S_2^- = \{y \in S_2 : y \leqslant y_j\}$, observing that $|S_1^-| = i_0$ and $|S_2^-| = j$. Since ℓ is of positive slope, it follows that, if $(x, y) \in \ell$ for any $x \in S_1^-$, we must have $y \in S_2^-$.

If $j \geqslant i_0$, then, since each value in S_1^- can correspond to at most one value in S_2^-, it follows that there will be at least $j - i_0$ coordinates in S_2^-, and thus in S_2, that are not mapped to by ℓ, and so ℓ can contain at most $n - (j - i_0)$ points of Γ. Similarly, if $j < i_0$, since each value in S_2^- corresponds to at most one value from S_1^-, there will be at least $i_0 - j$ values in S_1 not covered by ℓ, and so ℓ contains at most $n - (i_0 - j)$ points from Γ. This shows that lines of positive slope through $(0, y_j)$ can contain at most $n - |j - i_0|$ points of Γ.

For lines of negative slope, we can reflect the grid in the x-axis, considering $S_2' = \{-y : y \in S_2\}$. This reverses the ordering of the elements, so $y_i' = -y_{n-i+1}$. The line ℓ then corresponds to a line ℓ' of positive slope passing through $(0, y_{n-j+1}')$, and thus contains at most $n - |(n - j + 1) - i_0|$ points of the grid. $\qquad \square$

With this claim in mind, we now define an improved weighting of the points of Γ. As in the claim, enumerate the elements of S_1 as $x_1 < x_2 < \ldots < x_n$ and those of S_2 as $y_1 < y_2 < \ldots < y_n$, and let $i_0, j_0 \in [n]$ be such that $x_{i_0} = 0$ and $y_{j_0} = 0$. Given parameters α, β and t, to be chosen later, we define the weights on $\Gamma \setminus \{(0,0)\}$ as follows:

$$
w((x_i, y_j)) = \begin{cases} \alpha & \text{if } i \neq i_0 \text{ and } j \neq j_0; \\ \beta & \text{if } j = j_0, \text{ or if } i = i_0 \text{ and } \min\{|j - i_0|, |n - j + 1 - i_0|\} \geqslant t; \\ 0 & \text{otherwise.} \end{cases}
$$

That is, we assign weight α to all nonaxial points and weight β to all axial points except those in intervals around y_{i_0} and y_{n-i_0+1} on the y-axis.

For the weighting to be valid, we require that each origin-avoiding line have weight at most 1. If the line ℓ contains two axial points with positive weight, then let it pass through $(0, y_j)$ on the y-axis. By definition of w, we must have $|j - i_0|, |(n - j + 1) - i_0| \geqslant t$, and so by Claim 4.1, ℓ contains at most $n - t$ points of Γ, and hence at most $n - t - 2$ nonaxial points. Thus, the weight of any such line is at most $(n - t - 2)\alpha + 2\beta$.

Otherwise, the line ℓ can contain at most one weighted axial point and at most $n - 1$ nonaxial points, giving a total weight of not more than $(n - 1)\alpha + \beta$. Hence, our parameters must satisfy $(n - t - 2)\alpha + 2\beta \leqslant 1$, $(n - 1)\alpha + \beta \leqslant 1$, $\alpha, \beta \geqslant 0$, and $t \in \mathbb{N}$.

With regards to the objective function, we note that there are at most $2(2t - 1)$ axial points with weight zero, and thus $w(\Gamma) \geqslant (n - 1)^2\alpha + 2(n - 2t)\beta$, and we wish to maximise this quantity subject to the constraints above. Some routine calculations then yield that we should set $t = \left\lceil \frac{1}{2}\sqrt{(5n + 1)(n - 1)} - n \right\rceil$, $\alpha = \frac{1}{n+t}$ and $\beta = \frac{t+1}{n+t}$, for which we have $w(\Gamma) \geqslant \left(10 - 4\sqrt{5} + O(n^{-1})\right)(n - 1)$. Proposition 2.1 gives the desired lower bound:

$$\mathrm{cov}_k(\Gamma) \geqslant k\Phi(\Gamma) \geqslant kw(\Gamma) \geqslant \left(10 - 4\sqrt{5} + O(n^{-1})\right)k(n - 1).$$

(c) For the final part of the theorem, we assume the grid Γ is Δ-generic, meaning that any line through two axial points can contain at most Δ nonaxial points. In the framework of part (b), where we assign a weight of α to all nonaxial points, and a weight of β to all axial points, we then obtain the constraint $(n - 1) + \alpha + \beta \leqslant 1$ from the lines with at most one axial point, and the constraint $\Delta\alpha + 2\beta \leqslant 1$ from the lines with two axial points. The total weight, which we seek to maximise, is $(n - 1)^2\alpha + 2(n - 1)\beta$.

This optimisation problem is solved by taking $\alpha = \frac{1}{2(n-1)-\Delta}$ and $\beta = 1 - \frac{n-1}{2(n-1)-\Delta}$. It is readily verified that both constraints are then satisfied with equality, and the total weight of the grid is $\left[2 - \frac{n-1}{2(n-1)-\Delta}\right](n - 1)$, whence the result follows by once again applying Proposition 2.1. □

A few remarks are in order at this point. First, we note that the lower bound of part (b) can be improved if we have additional information about where the origin is in the grid. For example, suppose the origin is in the lower-left corner; that is, $\min S_1 = \min S_2 = 0$. Then any line containing two axial points must be of negative slope, and hence when we apply Claim 4.1, we see that it is enough to leave only the largest values on the y-axis unweighted. When one solves the corresponding optimisation problem, we find that $\mathrm{cov}_k(\Gamma) \geqslant \left(4 - 2\sqrt{2} + O(n^{-1})\right)k(n - 1)$ for such grids Γ, a considerable improvement in the constant factor, as $4 - 2\sqrt{2} \approx 1.1716$.

Second, as explained in the introduction, almost all grids Γ are generic, and parts (a) and (c) of Theorem 1.5 determine $\mathrm{cov}_k(\Gamma)$ precisely when k is even and asymptotically when k is odd and large. However, even when the grid Γ is not generic but only Δ-generic, provided $\Delta = o(n)$, part (c) is robust enough to resolve the problem asymptotically. We give some natural examples of this below.

Corollary 4.2. *Given $n \in \mathbb{N}$, let $\Gamma_{\mathrm{exp},n} = \Gamma(E, E)$, where $E = \{0, 1, 2, 4, \ldots, 2^{n-2}\}$, and let $\Gamma_{\mathrm{quad},n} = \Gamma(S, S)$, where $S = \{0, 1, 4, \ldots, (n - 1)^2\}$. Then, if $\Gamma \in \{\Gamma_{\mathrm{exp},n}, \Gamma_{\mathrm{quad},n}\}$, we have $\mathrm{cov}_k(\Gamma) = \left(\frac{3}{2} + o(1)\right)k(n - 1)$ as $k, n \to \infty$.*

Proof. By Theorem 1.5(a), we know $\mathrm{cov}_k(\Gamma) \leqslant \left\lceil \frac{3}{2}k \right\rceil (n-1)$, which is $\left(\frac{3}{2} + O(k^{-1}) \right) k(n-1)$, and so we need only demonstrate the lower bound. By Theorem 1.5(c), it suffices to show Γ is Δ-generic for some $\Delta = o(n)$.

We begin with $\Gamma = \Gamma_{\exp,n}$, and show that is 1-generic. Indeed, suppose ℓ is a line containing two axial points of Γ, say $(0, 2^i)$ and $(2^j, 0)$. It is then a straightforward calculation to see that, for any r, the line ℓ passes through $(2^r, 2^i - 2^{i-j+r})$. In order for this to be a nonaxial point of the grid $\Gamma_{\exp,n}$, we require that $2^i - 2^{i-j+r}$ is a positive power of 2 strictly smaller than 2^i, which only happens for $r = j - 1$. Thus, the line ℓ contains exactly one nonaxial point, and hence $\Gamma_{\exp,n}$ is 1-generic.

The quadratic grid $\Gamma_{\mathrm{quad},n}$ requires somewhat more delicate treatment. Again, let us suppose ℓ is a line containing the axial points $(0, a^2)$ and $(b^2, 0)$, and let (r^2, s^2) be a nonaxial point lying on ℓ. We then have $s^2 = a^2 \left(1 - \frac{r^2}{b^2} \right)$, or $(ab)^2 = (ar)^2 + (bs)^2$. Thus, we can bound the number of nonaxial points on ℓ by the number of ways of writing $(ab)^2$ as a sum of two positive squares. Following Beiler [Bei64, pp. 140–142], if Q is the set of prime divisors of ab that are congruent to 1 modulo 4, and if β_q is the multiplicity of $q \in Q$ in the prime factorisation of $(ab)^2$, then there are $\prod_{q \in Q}(\beta_q + 1) - 1$ ways to write $(ab)^2$ as a sum of squares.

To simplify the notation in the calculation below, we shall assume $q = 5, 13$ and 17 are included in Q, setting $\beta_q = 0$ in case they do not divide ab. Now observe that we have

$$\sum_{q \in Q}(\beta_q + 1) = 3 + \beta_5 + \beta_{13} + \beta_{17} + \sum_{q \in Q, q \geqslant 29} (\beta_q + 1)$$

$$\leqslant 3 + \beta_5 + \beta_{13} + \beta_{17} + \sum_{q \in Q, q \geqslant 29} 2\beta_q$$

$$\leqslant 3 + \sum_{q \in Q} \beta_q \log_5 q = 3 + \log_5 \left(\prod_{q \in Q} q^{\beta_q} \right) \leqslant 3 + \log_5 \left((ab)^2 \right),$$

and so $\sum_{q \in Q}(\beta_q + 1) \leqslant 3 + 4 \log_5 n$.

Since $m \leqslant 3^{m/3}$ for every $m \in \mathbb{N}$, we have $\prod_i m_i \leqslant 3^{\frac{1}{3} \sum_i m_i}$, and hence the number of ways to express $(ab)^2$ as a sum of squares is less than $3^{\frac{1}{3} \sum_{q \in Q}(\beta_q + 1)} \leqslant 3 \cdot 3^{\frac{4}{3} \log_5 n} = 3n^{\frac{4}{3} \log_5 3}$. It thus follows that $\Gamma_{\mathrm{quad},n}$ is Δ-generic for $\Delta = 3n^{\frac{4}{3} \log_5 3} = o(n)$. $\qquad \square$

4.2. Standard grids

While the results of Section 4.1 resolve the problem asymptotically for very many grids, there is no questioning the fact that the most natural case to consider is that of the standard grid $\Gamma(S, S)$, where $S = \{0, 1, \ldots, n-1\}$. For convenience, we denote this grid by Γ_n. By considering the line $x + y = n - 1$, we see that Γ_n is not Δ-generic for any $\Delta < n - 2$, which means the lower bound of Theorem 1.5(c) is worse than the Ball–Serra bound for $n > k$. By tailoring our methods for this specific grid, though, we will obtain much better bounds. We begin by showing that a strict k-cover of the standard grid requires far fewer lines than the upper bound given by Theorem 1.5, which, as shown in the previous subsection, is asymptotically tight for most grids.

Proof of Theorem 1.6 (upper bound). We prove the upper bound by constructing a small k-cover of Γ_n. Let $t \in [n-1]$ be a parameter, to be determined later, and consider the following collection of lines:

(i) $\left\lceil \frac{i}{n+t-1} k \right\rceil$ copies of the lines $x = i$ and $y = i$ for each $i \in [n-1]$;

(ii) $k - \left\lceil \frac{i}{n+t-1} k \right\rceil$ copies of the line $x + y = i$ for every $1 \leqslant i < n + t - 1$.

We begin by showing that the above collection of lines gives a k-cover of Γ_n. First, it is clear that no line passes through the origin. Now consider a point $(s_1, s_2) \in \Gamma_n \setminus \{(0,0)\}$. If $s_1 + s_2 < n + t - 1$, then (s_1, s_2) is covered $\left\lceil \frac{s_1}{n+t-1} k \right\rceil + \left\lceil \frac{s_2}{n+t-1} k \right\rceil$ times by the lines in (i) and another $k - \left\lceil \frac{s_1+s_2}{n+t-1} k \right\rceil$ times by those in (ii), and thus at least k times in total. On the other hand, if $s_1 + s_2 \geqslant n + t - 1$, then lines in (i) alone cover the point $\left\lceil \frac{s_1}{n+t-1} k \right\rceil + \left\lceil \frac{s_2}{n+t-1} k \right\rceil \geqslant k$ times, as required. Calculating the size of this k-cover, we obtain

$$
\begin{aligned}
\mathrm{cov}_k(\Gamma_n) &\leqslant 2 \sum_{i=1}^{n-1} \left\lceil \frac{i}{n+t-1} k \right\rceil + \sum_{i=1}^{n+t-2} \left(k - \left\lceil \frac{i}{n+t-1} k \right\rceil \right) \\
&\leqslant k \left[2 \sum_{i=1}^{n-1} \frac{i}{n+t-1} + \sum_{j=1}^{n+t-2} \frac{j}{n+t-1} \right] + 2n \\
&\leqslant k \left[\frac{2}{n+t-1} \binom{n}{2} + \frac{1}{n+t-1} \binom{n+t-1}{2} \right] + 2n \\
&= k \left[\frac{n(n-1)}{n+t-1} + \frac{n+t-2}{2} \right] + 2n.
\end{aligned}
\tag{4.1}
$$

The upper bound given by (4.1) is valid for any $t \in [n-1]$; we now want to choose a value of t that makes the right-hand side as small as possible. The function $g(t) = \frac{n(n-1)}{n+t-1} + \frac{n+t-2}{2}$ has its minimum at $t_0 = \sqrt{2(n-1)n} - (n-1)$. Since our parameter must be an integer, we choose $t = \left\lceil \sqrt{2(n-1)n} - (n-1) \right\rceil = (\sqrt{2} - 1 + O(n^{-1}))(n-1)$, for which we have $g(t) = (\sqrt{2} + O(n^{-1}))(n-1)$. Substituting this into (4.1) yields the claimed upper bound of $(\sqrt{2} + O(n^{-1} + k^{-1})) k(n-1)$ on $\mathrm{cov}_k(\Gamma_n)$. $\qquad\square$

We now turn our attention to the lower bound.

Proof of Theorem 1.6 (lower bound). By Proposition 2.1, it suffices to find a feasible solution to the dual linear program $\mathcal{D}(\Gamma_n)$, that is, a weighting of the nonzero points of Γ_n in which every origin-avoiding line has weight at most 1, that has total weight at least $(2 - e^{-1/2} + O(n^{-1}))(n-1)$.

Let t be the largest integer such that $\sum_{i=1}^{t} \frac{1}{n-i} \leqslant \frac{1}{2}$ and consider the following weighting on the points of $\Gamma_n \setminus \{(0,0)\}$:

$$
w((x,y)) = \begin{cases} \frac{1}{2} & \text{if } xy = 0; \\ \frac{1}{n-i} & \text{if } x + y = n - 1 + i \text{ for some } i \in [t]; \\ 0 & \text{otherwise.} \end{cases}
$$

We first show that $(w((x, y)) : (x, y) \in \Gamma_n)$ gives a feasible solution to the dual linear program $\mathcal{D}(\Gamma_n)$. Clearly $w((x, y)) \geqslant 0$ for all $(x, y) \in \Gamma_n \setminus \{(0, 0)\}$. Now, let $\ell \in \mathcal{L}$ be any line. If ℓ contains two axial points, then any nonaxial point (x, y) on ℓ satisfies $x + y \leqslant n - 1$, and thus has weight zero. It follows that $w(\ell) = 1$. Otherwise, if $(x, y) \mapsto x + y$ is constant on ℓ, then $\ell = \{(x, y) : x + y = n - 1 + i\}$ for some $i \in [n - 1]$. Then all points on ℓ have weight zero, unless $1 \leqslant i \leqslant t$, in which case ℓ contains $n - i$ points of weight $\frac{1}{n-i}$ each. Thus $w(\ell) \leqslant 1$ in this case. Finally, if $(x, y) \mapsto x + y$ is not constant on ℓ, it must be injective. Then, ℓ contains at most one axial point, which has weight $\frac{1}{2}$, and the weight from the remaining points is at most $\sum_{i=1}^{t} \frac{1}{n-i}$, which by the choice of t is at most $\frac{1}{2}$. So in total we again have $w(\ell) \leqslant 1$.

To compute the total weight of the grid, observe that each diagonal line of the form $x + y = i$ has weight one if $1 \leqslant i \leqslant n - 1 + t$ and zero otherwise. Thus, the total weight of the grid is $n - 1 + t$.

It remains to estimate t. Note that $t \leqslant \frac{n}{2}$, since $\sum_{i=1}^{t} \frac{1}{n-i} \geqslant \sum_{i=1}^{t} \frac{1}{n} = \frac{t}{n}$, and so both $n - 1$ and $n - 1 - t$ go to infinity linearly with n. It is well known that, as $m \to \infty$, the partial sums H_m of the Harmonic series satisfy $H_m = \sum_{j=1}^{m} \frac{1}{j} = \log m + \gamma + O(m^{-1})$, where γ is a constant. Hence,

$$\sum_{i=1}^{t} \frac{1}{n-i} = H_{n-1} - H_{n-1-t} = \log\left(\frac{n-1}{n-1-t}\right) + O(n^{-1}).$$

Thus, we must have $\log\left(\frac{n-1}{n-1-t}\right) = \frac{1}{2} + O(n^{-1})$, or $\log\left(1 - \frac{t}{n-1}\right) = -\frac{1}{2} + O(n^{-1})$. This gives $1 - \frac{t}{n-1} = e^{-1/2} + O(n^{-1})$, or $t = (1 - e^{-1/2} + O(n^{-1}))(n - 1)$, which results in the claimed bound. $\qquad\square$

The cover we constructed for the upper bound only uses lines of slope 0, ∞, and -1, which we shall hereon refer to as horizontal, vertical, and diagonal lines respectively. It may seem rather limiting to restrict ourselves to these lines, and it is natural to wonder if one can do better by using other lines as well. However, we verified by computer search that for small values of n and k, one is always able to build an optimal k-cover consisting only of these three types of lines. This motivated us to further study this restricted class of k-covers, and in our final result we prove that the smallest k-cover of this form has size at least $\left(\sqrt{2} + O(n^{-1})\right)k(n - 1)$.

To improve the lower bound in this setting, we again appeal to the linear programming approach and assign weights to the points of the grid. Since we are only permitted to use horizontal, vertical, and diagonal lines, the dual linear program only has constraints on the total weights of these lines. This gives us much more freedom in choosing the weights, and so we can hope to find a feasible weighting with a larger sum, matching the upper bound from our construction.

Indeed, we now have $n^2 - 1$ degrees of freedom (the weights of the individual nonzero points) but only $4(n - 1)$ constraints (the weights of the horizontal, vertical, and diagonal lines). In order to make the problem analytically feasible, we will reduce the search space by imposing some further restrictions on the weighting.

First, we observe that the principle of complementary slackness dictates that for any line that appears in an optimal cover, the corresponding constraint in the dual problem must be satisfied with equality. In particular, since our construction uses all horizontal and vertical lines, together

with the diagonal lines $x + y = i$ when i is not too large, we will ensure that the weights of these lines are exactly 1.

Second, to reduce the number of variables from quadratic to linear, we will require that the weighting is constant along the nonaxial points on diagonal lines. As we shall show below, these additional requirements allow us to set up and solve a recurrence relation for the weights, which yields the desired result.

Proof of Theorem 1.7. In light of the discussion above, we define, for some parameters $t \in \mathbb{N}$ and $\beta_1, \beta_2, \ldots, \beta_{n-1}, z \in \mathbb{R}_{\geqslant 0}$, the following weighting of the grid:

$$w((x, y)) = \begin{cases} \frac{1}{2} - \beta_{x+y} & \text{if } x = 0 \text{ or } y = 0; \\ \frac{2\beta_{x+y}}{x+y-1} & \text{if } x, y \neq 0 \text{ and } 1 \leqslant x + y \leqslant n - 1; \\ \frac{1}{2n-1-i} & \text{if } x + y = i \text{ for some } n \leqslant i \leqslant n + t - 1; \\ z & \text{if } x + y = n + t; \\ 0 & \text{if } x + y \geqslant n + t + 1. \end{cases}$$

We start by inspecting the diagonal lines $x + y = i$. If $1 \leqslant i \leqslant n - 1$, the line has two axial points of weight $\frac{1}{2} - \beta_i$ each, and $i - 1$ nonaxial points weighing $\frac{2\beta_i}{i-1}$ each. Thus, the total weight of the line is 1. If instead we have $n \leqslant i \leqslant n + t - 1$, then the diagonal line consists of $2n - 1 - i$ nonaxial points, which each have a weight of $\frac{1}{2n-1-i}$, again giving a total weight of 1. When $i = n + t$, the line consists of $n - t - 1$ points of weight z, and so feasibility dictates that $0 \leqslant z \leqslant \frac{1}{n-t-1}$. Finally, if $i \geqslant n + t + 1$, the line has weight 0.

By summing along these diagonal lines, we find that $n + t - 1 \leqslant w(\Gamma_n \setminus \{(0,0)\}) \leqslant n + t$, and so our goal is to maximise t while ensuring that the weighting is feasible. We are yet to specify the values of the parameters β_i, which we will deduce from the requirement that the vertical lines all have weight 1. Note that the weighting is symmetric in the line $y = x$, and so this will also ensure that all horizontal lines have weight 1.

Consider the line $x = n - 1$, which has weight $\frac{1}{2} - \beta_{n-1} + \sum_{i=n}^{n+t-1} \frac{1}{2n-1-i} + z$. Setting this equal to 1 yields

$$\beta_{n-1} = \sum_{i=n}^{n+t-1} \frac{1}{2n-1-i} + z - \frac{1}{2}. \tag{4.2}$$

Now compare the weights of the points on the lines $x = n - 2$ and $x = n - 1$. Since the diagonals are constant along their nonaxial points, we have $w(n - 2, y) = w(n - 1, y - 1)$ for all $2 \leqslant y \leqslant n - 1$. Hence the differences are that $w(n - 2, 0)$ and $w(n - 2, 1)$ replace $w(n - 1, 0)$ and $w(n - 1, n - 1)$. Thus,

$$\begin{aligned} w(\{x = n - 2\}) &- w(\{x = n - 1\}) \\ &= w(n - 2, 0) + w(n - 2, 1) - w(n - 1, 0) - w(n - 1, n - 1) \\ &= \tfrac{1}{2} - \beta_{n-2} + \frac{2\beta_{n-1}}{n - 2} - \left(\tfrac{1}{2} - \beta_{n-1}\right) - 0, \end{aligned}$$

and since both vertical lines have weight 1, this gives $\beta_{n-2} = \left(1 + \frac{2}{n-2}\right)\beta_{n-1}$. Repeating this argument for the lines $x = i - 1$ and $x = i$ for each $2 \leqslant i \leqslant n - 1$, we obtain the following recurrence relation:

$$\beta_{i-1} = \left(1 + \frac{2}{i-1}\right)\beta_i - z\mathbf{1}_{i=t+1} - \frac{1}{n-i}\mathbf{1}_{i\leqslant t}, \tag{4.3}$$

where $\mathbf{1}_A$ is the indicator function of the event A defined as

$$\mathbf{1}_A = \begin{cases} 1 & \text{if } A \text{ is true;} \\ 0 & \text{if } A \text{ is false.} \end{cases}$$

For the initial condition, observe that the line $x + y = 1$ has no nonaxial points, and so for it to have weight 1, we must have $\beta_1 = 0$. Combining this with (4.3), we obtain:

$$\beta_i = \frac{1}{i(i+1)} \sum_{j=1}^{\min\{t,i\}} \frac{j(j-1)}{n-j} + \mathbf{1}_{i\geqslant t+1}\frac{t(t+1)}{i(i+1)}z \quad \text{for all } 2 \leqslant i \leqslant n-1. \tag{4.4}$$

Recall that we require $0 \leqslant z \leqslant \frac{1}{n-t-1}$. Substituting the value of β_{n-1} from (4.4) into (4.2), we can solve for z to obtain:

$$z = \left(\frac{1}{2} - \sum_{j=1}^{t} \frac{1}{n-j}\left(1 - \frac{j(j-1)}{(n-1)n}\right)\right)\frac{n(n-1)}{n(n-1)-t(t+1)}$$

$$= \frac{(n-1)n\left(\frac{1}{2} - \frac{t(2n+t-1)}{2(n-1)n}\right)}{(n-1)n - t(t+1)},$$

where the second equality is due to the fact that

$$\sum_{j=1}^{t} \frac{1}{n-j}\left(1 - \frac{j(j-1)}{(n-1)n}\right) = \frac{1}{n(n-1)} \sum_{j=1}^{t}(n+j-1) = \frac{t(2n+t-1)}{2n(n-1)}.$$

If $n > 1$, we have $z \geqslant 0$ when $0 \leqslant t \leqslant \frac{1}{2}\left(\sqrt{8n^2 - 8n + 1} - 2n + 1\right)$ and $z \leqslant \frac{1}{n-t-1}$ for $\frac{1}{2}\left(\sqrt{8n^2 - 8n + 1} - 2n - 1\right) \leqslant t < n - 1$. These inequalities are satisfied by taking $t = \left\lfloor\frac{1}{2}(\sqrt{8n^2 - 8n + 1} - 2n + 1)\right\rfloor$, for which we have $t = \left(\sqrt{2} - 1 + O(n^{-1})\right)(n-1)$. It follows that the total weight of the grid is $\left(\sqrt{2} + O(n^{-1})\right)(n-1)$.

We are not quite done, as there is one final condition to verify — to ensure that all our weights are nonnegative, we must have $0 \leqslant \beta_i \leqslant \frac{1}{2}$ for all $1 \leqslant i \leqslant n-1$. From (4.3) we have:

$$\beta_i = \frac{i-1}{i+1}\left(\beta_{i-1} + \frac{1}{n-i}\right) \qquad\qquad\qquad \text{if } 2 \leqslant i \leqslant t.$$

$$\beta_i = \frac{i-1}{i+1}(\beta_{i-1} + z) \leqslant \frac{i-1}{i+1}\left(\beta_{i-1} + \frac{1}{n-i}\right) \qquad\qquad \text{if } i = t+1$$

$$\beta_i = \frac{i-1}{i+1}\beta_{i-1} < \beta_{i-1} \qquad\qquad\qquad\qquad \text{if } t+2 \leqslant i \leqslant n-1$$

Thus, it suffices to show that $\beta_i \leqslant \frac{1}{2}$ for all $2 \leqslant i \leqslant t + 1$. We will do so by showing that for $1 \leqslant i \leqslant t + 1$, we have $\beta_i \leqslant \frac{i-1}{2(n-i-1)}$ by induction on i. We know that $\beta_1 = 0$, so the base case is clear. Let $i > 1$ and assume the induction hypothesis; then

$$\beta_i \leqslant \frac{i-1}{i+1}\left(\beta_{i-1} + \frac{1}{n-i}\right) \leqslant \frac{i-1}{i+1}\left(\frac{i-2}{2(n-i)} + \frac{1}{n-i}\right) \leqslant \frac{i-1}{2(n-i-1)}.$$

We have $\frac{i-1}{2(n-i-1)} \leqslant \frac{1}{2}$ whenever $i \leqslant n - i$, which is true since $i \leqslant t + 1 = (\sqrt{2} - 1 + O(n^{-1}))(n-1)$.

Hence our weighting is indeed feasible and $w(\Gamma_n \setminus \{(0,0)\}) = (\sqrt{2} + O(n^{-1}))(n-1)$. It thus follows that any k-cover of Γ_n using only horizontal, vertical, and diagonal lines must have size at least $(\sqrt{2} + O(n^{-1}))k(n-1)$. $\qquad\square$

5. Conclusion

In this paper, we studied line coverings with multiplicities for two-dimensional real grids. We determined the minimum size of a cover in several cases, but some natural and interesting questions remain open, and we highlight them below.

In Section 3, we investigated for which grids the Ball–Serra bound is tight. We proved that, when n is sufficiently large with respect to m and k, the Ball–Serra bound is tight for any $n \times m$ grid. Moreover, we showed that the threshold value for n given by Theorem 1.2 is tight for *most* grids. It can be shown, however, that this bound on n is not best possible for *all* grids. For example, for the grid $\Gamma(S_1, S_2)$, where $S_1 = \{0, 1, 2, \ldots, n-1\}$ and $S_2 = \{-1, 0, 1\}$, and any $k \geqslant 3$, we can show that the Ball–Serra bound is tight already for $n = 2(k-1)$, as opposed to the lower bound $n \geqslant 2k - 1$ of Theorem 1.2. However, in this grid the omitted point $(0,0)$ is not a corner point, while, as in the square grid setting, it is more natural to consider grids in which $(0,0)$ is a corner. For such grids, one could investigate when the Ball–Serra bound holds.

Question 5.1. Let Γ be the grid $\{0, 1, 2, \ldots, n-1\} \times \{0, 1, 2, \ldots, m-1\}$ and $k \geqslant 2$ be an integer. How large must n be with respect to m and k to have $\mathrm{cov}_k(\Gamma) = k(n-1) + (m-1)$?

Our main result for standard grids establishes reasonably good asymptotic lower and upper bounds on $\mathrm{cov}_k(\Gamma_n)$. It would be of interest to close the remaining gap.

Question 5.2. What is the true asymptotic value of $\mathrm{cov}_k(\Gamma_n)$?

We tend to believe that $\mathrm{cov}_k(\Gamma_n) = (\sqrt{2} + o(1))k(n-1)$. In Theorem 1.7, we showed this to be the case when we only use lines of slope 0, ∞, and -1. However, for the weighting we used to establish the lower bound, one can show that lines of slope 1 near the origin (e.g., $y = x + 1$) have weight larger than 1 when n is large. We believe that these are the only problematic lines, and so as an intermediate step one could attempt to verify that our weighting remains feasible if one only forbids lines of slope 1. This would imply that any k-cover of Γ_n of size smaller than $(\sqrt{2} + o(1))k(n-1)$ must contain many lines of slope 1. To show that such a construction is unlikely to exist, it might be helpful to consider what happens if we restrict ourselves to lines of slope 0, ∞, -1, and 1.

In our work thus far we observed that the standard grid Γ_n requires many fewer lines to cover than any other $n \times n$ grid we considered. Our general lower bound from Theorem 1.5(b) (and the improvement for grids in which $(0,0)$ is a corner discussed after the proof) is not strong enough to establish this fact, and we propose the following problem.

Question 5.3. Is it true that $\text{cov}_k(\Gamma_n) \leqslant \text{cov}_k(\Gamma)$ for any $n \times n$ grid Γ in which $(0,0)$ is a corner?

More broadly, it would be of interest to improve the lower bound from Theorem 1.5(b), which we do not believe to be best possible.

Another direction that might lead to interesting findings is to consider translates of the standard grid in which the omitted point is not in the lower-left corner. How does the position of the origin then affect the value of $\text{cov}_k(\Gamma)$? For instance, what are the asymptotics of $\text{cov}_k(\Gamma)$, where $\Gamma = \Gamma(\{-\lfloor n/2 \rfloor, \ldots, \lceil n/2 \rceil\}, \{-\lfloor n/2 \rfloor, \ldots, \lceil n/2 \rceil\})$?

While we mainly focused on two-dimensional real grids, it would be natural to investigate the problem in higher dimensions as well. For example, in Section 3, we remarked that the Ball–Serra bound can be tight for higher-dimensional grids as well, provided that one of the sides is much longer than the others. How much longer does that side need to be for the bound to be attained? Once again, it would be particularly interesting to investigate $\text{cov}_k(\Gamma_n^{(d)})$ for the standard d-dimensional grid $\Gamma_n^{(d)} = \{0, \ldots, n-1\}^d$.

Finally, while all of our results are stated for grids over $\mathbb{R}$, the questions we considered and our general framework extend to grids over any field. Some of our results, for example Theorem 1.2, extend to arbitrary fields. It will be interesting to prove similar results for other fields, and in particular for fields of positive characteristic.

Acknowledgements

References

[ABNR22] Gianira N. Alfarano, Martino Borello, Alessandro Neri, and Alberto Ravagnani. Three combinatorial perspectives on minimal codes. *SIAM J. Discrete Math.*, 36(1):461–489, 2022. `doi:10.1137/21M1391493`.

[AF93] Noga Alon and Zoltán Füredi. Covering the cube by affine hyperplanes. *European J. Combin.*, 14(2):79–83, 1993. `doi:10.1006/eujc.1993.1011`.

[Alo99] Noga Alon. Combinatorial Nullstellensatz. *Combin. Probab. Comput.*, 8(1-2):7–29, 1999. `doi:10.1017/S0963548398003411`.

[BBDM23] Anurag Bishnoi, Simona Boyadzhiyska, Shagnik Das, and Tamás Mészáros. Subspace coverings with multiplicities. *Combin. Probab. Comput.*, 32(5):782–795, 2023. `doi:10.1017/s0963548323000123`.

[BBS10]   Aart Blokhuis, Andries E. Brouwer, and Tamás Szőnyi. Covering all points except one. *J. Algebraic Combin.*, 32(1):59–66, 2010. `doi:10.1007/s10801-009-0204-1`.

[BCPS18]   Anurag Bishnoi, Pete L. Clark, Aditya Potukuchi, and John R. Schmitt. On zeros of a polynomial in a finite grid. *Combin. Probab. Comput.*, 27(3):310–333, 2018. `doi:10.1017/S0963548317000566`.

[Bei64]   Albert H. Beiler. *Recreations in the Theory of Numbers: The Queen of Mathematics Entertains*. Dover Publications, New York, 1964.

[BGG⁺22]   Bartłomiej Bosek, Jarosław Grytczuk, Grzegorz Gutowski, Oriol Serra, and Mariusz Zajac. Graph polynomials and group coloring of graphs. *European J. Combin.*, 102:Paper No. 103505, 11, 2022. `doi:10.1016/j.ejc.2021.103505`.

[Bru92]   Aiden A. Bruen. Polynomial multiplicities over finite fields and intersection sets. *J. Combin. Theory Ser. A*, 60(1):19–33, 1992. `doi:10.1016/0097-3165(92)90035-S`.

[BS09]   Simeon Ball and Oriol Serra. Punctured combinatorial Nullstellensätze. *Combinatorica*, 29(5):511–522, 2009. `doi:10.1007/s00493-009-2509-z`.

[CH20]   Alexander Clifton and Hao Huang. On almost k-covers of hypercubes. *Combinatorica*, 40(4):511–526, 2020. `doi:10.1007/s00493-019-4221-y`.

[dB21]   Yvonne den Bakker. *Hyperplane covering problems*. Bachelor's thesis, Technische Universiteit Delft, 2021.

[GMCS22]   Stefan Glock, David Munhá Correia, and Benny Sudakov. The n-queens completion problem. *Res. Math. Sci.*, 9(3):Paper No. 41, 20, 2022. `doi:10.1007/s40687-022-00335-1`.

[Gut16]   Larry Guth. *Polynomial methods in combinatorics*, volume 64 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2016. `doi:10.1090/ulect/064`.

[Jam77]   Robert E. Jamison. Covering finite fields with cosets of subspaces. *J. Combinatorial Theory Ser. A*, 22(3):253–266, 1977. `doi:10.1016/0097-3165(77)90001-2`.

[Kom94]   Péter Komjáth. Partitions of vector spaces. *Period. Math. Hungar.*, 28(3):187–193, 1994. `doi:10.1007/BF01876338`.

[LR14]   Ivan Landjev and Assia Rousseva. On the sharpness of Bruen's bound for intersection sets in Desarguesian affine spaces. *Des. Codes Cryptogr.*, 72(3):551–558, 2014. `doi:10.1007/s10623-012-9783-2`.

[MG07]   Jiří Matoušek and Bernd Gärtner. *Understanding and using linear programming*. Springer Berlin, Heidelberg, 2007. `doi:10.1007/978-3-540-30717-4`.

[SW22]   Lisa Sauermann and Yuval Wigderson. Polynomials that vanish to high order on most of the hypercube. *J. Lond. Math. Soc. (2)*, 106(3):2379–2402, 2022. `doi:10.1112/jlms.12637`.

[Zan02]   Corrado Zanella. Intersection sets in $AG(n, q)$ and a characterization of the hyperbolic quadric in $PG(3, q)$. *Discrete Math.*, 255(1-3):381–386, 2002. Combinatorics '98 (Palermo). `doi:10.1016/S0012-365X(01)00413-7`.