Near-linear-time algorithm for the geodetic Radon number of grids[☆]

Mitre Costa Dourado[a], Vinícius Gusmão Pereira de Sá[a,*], Dieter Rautenbach[b], Jayme Luiz Szwarcfiter[a]

[a] *Instituto de Matemática, Universidade Federal do Rio de Janeiro, Brazil*
[b] *Institut für Optimierung und Operations Research, Universität Ulm, Germany*

Abstract

The Radon number of a graph is the minimum integer r such that all sets of at least r of its vertices can be partitioned into two subsets whose convex hulls intersect. Determining the Radon number of general graphs in the geodetic convexity is NP-hard. In this paper, we show the problem is polynomial for d-dimensional grids, for all $d \geq 1$. The proposed algorithm runs in near-linear $\mathcal{O}(d\ (\log d)^{1/2})$ time for grids of arbitrary sizes, and in sub-linear $\mathcal{O}(\log d)$ time when all grid dimensions have the same size.

Keywords: graph convexity, geodetic convexity, Radon number, grids

---

1. Introduction

The concept of convexity in graphs was borrowed from its most well-known geometric counterpart, where a subset S of the Euclidean space is dubbed convex if, for every two points $x, y \in S$, the interval consisting of the straight segment connecting x and y is entirely contained in S. Formally, a *convexity space* (or simply *convexity*, for short) ϕ consists of a pair $(V, \mathcal{C})$, where V is a set—the *ground set*—and $\mathcal{C}$ is a collection of subsets of V—the *convex sets*—such that $\mathcal{C}$ contains both V and the empty set, and $\mathcal{C}$ is closed under arbitrary intersections and nested unions.

The ground set of a geometric convexity is a set of points. Analogously, the ground set of a graph convexity is the set of vertices of some connected

[☆] Work partially supported by CNPq and FAPERJ.
[*] Corresponding author. Phone: +55 (21) 3938-9955. E-mail: vigusmao@dcc.ufrj.br.

graph G. Many different geometric convexities have been studied to date (see, for instance, [3]) and the interval defined by two points is not always the straight segment between them in Euclidean fashion. Likewise, several distinct types of graph convexities have been considered in the literature [8], with applications ranging from statistical physics and distributed computing to marketing and social networks. In the *geodetic* graph convexity—also known as *geodesic* convexity [6]—a set $S \subseteq V(G)$ is convex if, for all $x, y \in S$, every *shortest path* between x and y in G is entirely contained in S.

Given a subset V' of a ground set V, the *convex hull* of V', denoted $[V']$, is the unique minimal convex subset of V containing V'. In the early 1920's, Johann Radon formulated a celebrated theorem stating that every set with at least $d+2$ points in $\mathbb{R}^d$ can be partitioned into two subsets whose convex hulls intersect [7]. A natural question concerns what happens when we consider some general ground set V instead of $\mathbb{R}^d$, and the *Radon number* of V is defined as the minimum integer r such that every subset of V with at least r elements can be partitioned in two sets whose convex hulls intersect. The Radon number of graphs has been used to model some problems occurring, for instance, in social networks.

A simple reduction from the maximum clique problem can be used to prove the NP-hardness of finding the Radon number of a graph in the geodetic convexity, hence a natural task is to determine such parameter for particular graph classes. We are interested in the class of d-dimensional grids, i.e., the Cartesian products of d paths of arbitrary sizes. A lot of insight on the problem was gained in [1]. In that paper, the authors derived general bounds and solved the problem for special cases. Computer-assisted results also disclosed (by brute force) the Radon number of all grids up to the ninth dimension. However, a polynomial-time algorithm for determining the Radon number of general grids was still outstanding.

In this paper, we introduce one such algorithm that runs in $\mathcal{O}(d\,(\log d)^{1/2})$ time. Additionally, a variation of our algorithm runs in sub-linear $\mathcal{O}(\log d)$ time for grids $G = P_n^d$, that is, those in which all d dimensions have the same size n.

2. The basics

If R is a subset of vertices of a graph G, then a partition $R = R_1 \cup R_2$ is a *Radon partition* if $[R_1] \cap [R_2] \neq \emptyset$. A set which admits no Radon partitions is called an *anti-Radon set*, also called a *Radon-independent set* by some authors. The Radon number of a graph G is thus the size of the maximum anti-Radon set of G plus one.

A grid $G = \text{Grid}(n_1, \ldots, n_d)$ is the Cartesian product of d paths $P_{n_1} \times P_{n_2} \times \ldots \times P_{n_d}$. The geodetic convexity on d-dimensional grids bears a natural resemblance with the convexity defined on the Euclidean space $\mathbb{R}^d$ by the Manhattan metric $(u, v) \mapsto ||u - v||_1$. If $R = \{u^1, \ldots, u^r\}$ is a set of vertices of $\text{Grid}(n_1, \ldots, n_d)$ with $u^i = (u^i_1, \ldots, u^i_d)$ for $i \in [1, r] := \{i \in \mathbb{N} : 1 \leq i \leq r\}$, then it is easy to see that the convex hull $[R]$ is the set of integer points in

$$\left[\min_{i\in[1,r]} u^i_1, \max_{i\in[1,r]} u^i_1\right] \times \left[\min_{i\in[1,r]} u^i_2, \max_{i\in[1,r]} u^i_2\right] \times \cdots \times \left[\min_{i\in[1,r]} u^i_d, \max_{i\in[1,r]} u^i_d\right]. \quad (1)$$

In other words, $[R]$ equals the Cartesian product of the one-dimensional convex hulls of the projections of R onto the d dimensions.

Having observed that, one can check whether a partition $R = R_1 \cup R_2$ is a Radon partition by simply inspecting the projections of R onto each dimension. If, for some $j \in [1, d]$, the greatest (smallest) coordinate of the projection of R_1 onto dimension ρ_j is less (greater) than the smallest (greatest) coordinate of the projection of R_2 on ρ_j, then the one-dimensional convex hulls of the projections of R_1 and R_2 onto ρ_j do not intersect, and $[R_1] \cap [R_2] = \emptyset$. In this case, we say the projections of R_1 onto ρ_j appear *all strictly to the left* (*all strictly to the right*) of the projections of R_2. If there is no such j, then the convex hulls of R_1 and R_2 intersect, and R is a Radon partition. Figure 1 illustrates the idea.

The problem of determining the Radon number of a grid G looks much harder. Indeed, not only is the number of partitions of a given set R exponential in $|R|$, but also the number of subsets R of the ground set $V(G)$ that would have to be checked in the worst case is exponential in $|V(G)|$.

In [2], Eckhoff determined the Radon number of the convexity space defined on $\mathbb{R}^d$ by the Manhattan metric $(u, v) \mapsto ||u - v||_1$ as

$$r(d) := \min\left\{r \in \mathbb{N} : \binom{r}{\lfloor \frac{r}{2} \rfloor} > 2d\right\}. \quad (2)$$

In [4], Jamison-Waldner observed that Eckhoff's result could be instantly leveraged to $\text{Grid}(n_1, \ldots, n_d)$ provided $n_j \geq r(d) - 1$, for all $j \in [1, d]$. However, if the grid dimensions are not as large, Eckhoff's result gives an upper bound (which may not be tight). In the next section, we obtain the exact geodetic Radon number of grids. We exploit the following theorem, rewritten

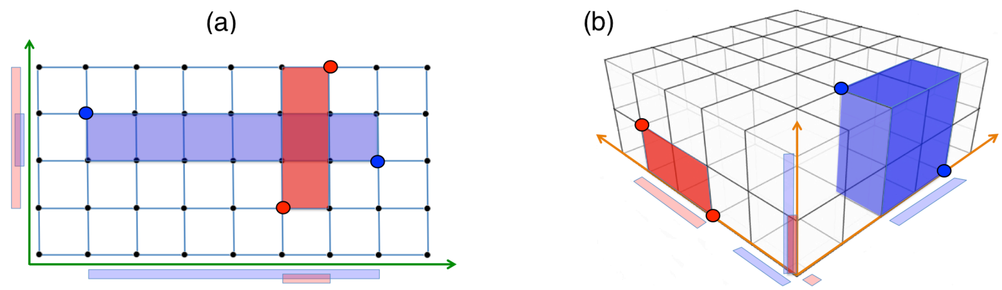

Figure 1: (a) Radon partition of a set containing four vertices of a 2-dimensional grid: on both dimensions the convex hulls of the projections of the partite sets intersect; (b) *not* a Radon partition of a set containing four vertices of a 3-dimensional grid: there are dimensions in which the projections of one of the partite sets appear all strictly to the left of the projections of the other partite set.

from [1], which characterizes grids with anti-Radon sets of size r in a suitable manner.

An *ordered partition* of a set V is a tuple $\psi = (V^1, \ldots, V^n)$, where $V^1 \cup \cdots \cup V^n$ is a partition of V.

Theorem 1 ([1]). *Let $d, n_1, \ldots, n_d, r \in \mathbb{N}$. The graph* $\mathrm{Grid}(n_1, n_2, \ldots, n_d)$ *has an anti-Radon set R with r vertices if and only if the set $[1, r]$ admits d ordered partitions $(V_1^1, \ldots, V_1^{n_1}), (V_2^1, \ldots, V_2^{n_2}), \ldots, (V_d^1, \ldots, V_d^{n_d})$ such that, for every subset S of $[1, r]$ with $1 \leq |S| \leq r/2$, there are indices $j \in [1, d]$ and $\ell \in [1, n_j]$ satisfying either $S = V_j^1 \cup \cdots \cup V_j^\ell$ or $S = V_j^{\ell+1} \cup V_j^{\ell+2} \cup \cdots \cup V_j^{n_j}$.*

The spirit of Theorem 1 is that the actual coordinates of a given set $R \subseteq V(G)$ do not really matter for the sake of deciding whether R is an anti-Radon set. What does matter is solely the sequence of orthogonal projections of R onto each dimension of the grid. As a matter of fact, because there may be coincident projections, what matters is the set of ordered partitions of R determined by its projections along each dimension (each partite set contains coinciding projections).

3. The algorithm

Let $G = \mathrm{Grid}(n_1, \ldots, n_d)$. The proposed algorithm checks whether $V(G)$ contains an anti-Radon set of size r, starting from Jamison-Waldner's upper bound $r = r(d) - 1$ and iteratively decrementing it, possibly all the way until $r = 2$.

We regard a non-trivial partition $R_1 \cup R_2$ of $[1, r]$, with $|R_1| = k \in [1, \lfloor r/2 \rfloor]$ without loss of generality, as a Radon partition *template* for any set $R \subseteq V(G)$ where $|R| = r$. We say it is a template of size k (or a k-*template*) for r vertices, and we let $\langle R_1 \rangle$ denote it. A template $\langle R_1 \rangle$ is *eliminated* by ordered partition ψ_j if all partite sets of ψ_j containing elements of R_1 appear either before the first partite set containing an element of $R_2 = [1, r] \setminus R_1$ or after the last partite set containing an element of R_2 (analogously, if $R_1 \cup R_2$ is a partition of some $R \subseteq V(G)$, the projections of R_1 appear either all strictly to the left or all strictly to the right of the projections of R_2 in dimension ρ_j).

By Theorem 1, it suffices to check, at any given iteration, if there exist d ordered partitions $\psi_1, \ldots, \psi_d$ of $[1, r]$ (the number of partite sets in each ψ_j matching the size n_j of the j-th dimension of G) such that template $\langle R_1 \rangle$ is eliminated, for each non-empty subset R_1 of $[1, r]$ with $|R_1| \leq r/2$. If that is the case, then G admits an anti-Radon set of size r. This way we avoid the burden of testing the anti-Radonness of an overall exponential number of subsets R of $V(G)$ in the worst case, even though the number of possible choices of d ordered partitions of $[1, r]$, for each r, is still obviously exponential.

We formulate three preliminary observations that set the basis for the algorithm to come. The *potential* of dimension ρ_j (with respect to a set of r vertices) stands for the maximum number of Radon partition templates for r vertices that can be eliminated by the j-th ordered partition ψ_j. Finally, the k-*quota* of dimension ρ_j with respect to a set of r vertices is the maximum number of k-templates that can be eliminated by ψ_j.

Observation 2 ([1]). *Given a subset R comprising r vertices of* Grid$(n_1, n_2, \ldots, n_d)$ *and some $j \in [1, d]$, the potential of dimension ρ_j is given by*

$$\textit{potential}_r(j) := \min\{n_j, r\} - 1.$$

Observation 3. *Given a set R with r elements and an integer $k \in [1, \lfloor r/2 \rfloor]$, the total number of k-templates for set R is given by*

$$\textit{template\_count}_r(k) := \begin{cases} \binom{r}{k} \cdot \frac{1}{2}, & \textit{if } k = \frac{r}{2} \\ \binom{r}{k}, & \textit{otherwise.} \end{cases}$$

Proof. If $k = r/2$, then each k-template comprises two complementary subsets of size k, hence the number of k-templates is half the number of subsets of size k. For other values of k, there is a bijection between the k-templates and the subsets of size k. □

Observation 4. *Given a set R with r elements and an integer $k \in [1, \lfloor r/2 \rfloor]$, the k-quota of dimension ρ_j with respect to R is given by*

$$k\text{-}quota_r(j) := \begin{cases} 1, & \text{if } k = \frac{r}{2} \text{ or } n_j = 2 \\ 2, & \text{otherwise.} \end{cases}$$

Proof. If $k \neq r/2$, the dimension ρ_j (analogously, the ordered partition ψ_j) may eliminate at most two k-templates: the k-template whose projections onto ρ_j appear all strictly to the left of the remaining projections, and the k-template whose projections appear all strictly to the right of the remaining projections, provided the size n_j of that dimension is greater than 2. If $n_j = 2$, though, then only one k-template may be eliminated, since projecting exactly k vertices onto the first coordinate of ρ_j and exactly k vertices onto the second coordinate of ρ_j would leave no room for the other $r - 2k > 0$ projections onto ρ_j.

If $k = r/2$, then two disjoint subsets R_1, R'_1 of $[1, r]$ with $|R_1| = |R'_1| = k$ are complementary, therefore any two k-templates $\langle R_1 \rangle, \langle R'_1 \rangle$ eliminated by an ordered partition ψ_j are actually one and the same bipartition of $[1, r]$. □

The pseudocode of the algorithm is shown as Algorithm 1. Each iteration of the main, outer loop of the algorithm consists of an attempt to prove the existence of an anti-Radon set of size r. It succeeds in doing so if it manages to prove the existence (without actually exhibiting them, for performance reasons) of d ordered partitions of $[1, r]$, under the size constraints imposed by the dimensions of grid G, so that all Radon partition templates for r vertices are eliminated.

For each r, the algorithm initializes two auxiliary structures (lines 2–3) which keep track of the quantities required by line 6. Then the inner loop of the algorithm (lines 4–9) considers each template size k, one at a time, from $\lfloor r/2 \rfloor$ to 1. The number of k-templates that must be eliminated is given by $template\_count_r(k)$ in Observation 3. Such k-templates are regarded as identical balls of color k that must be distributed into distinct bins (the dimensions), and the algorithm proceeds to distributing them (lines 5–9) in such a way that

(i) the total number of balls assigned to a bin never exceeds the potential of the corresponding dimension; and

Algorithm 1 Geodetic Radon number of grids

input: the dimension sizes n_j of a grid G, for $j = 1, \dots, d$
output: the Radon number of G

1 **for** $r = r(d) - 1, r(d) - 2, \dots, 2$ **do**
2 $balls[j] \leftarrow 0$, **for all** $j \in [1, d]$
3 $balls\_by\_color[j, k] \leftarrow 0$ **for all** $j \in [1, d], k \in [1, \lfloor r/2 \rfloor]$
4   **for** $k = \lfloor r/2 \rfloor, \lfloor r/2 \rfloor - 1, \dots, 1$ **do**
5   **repeat** $template\_count_r(k)$ **times** (see Observation 3)
6   find j such that $potential_r[j] - balls[j]$ is maximum,
              satisfying:
              (i) $balls[j] < potential_r[j]$ (see Observation 2)
              (ii) $balls\_by\_color[j, k] < k\text{-}quota_r(j)$ (see Observation 4)
7   **if** no such j exists,
              continue with the next r in line 1
8 $balls[j] \leftarrow balls[j] + 1$
9 $balls\_by\_color[j, k] \leftarrow balls\_by\_color[j, k] + 1$
10   **return** $r + 1$
11 **return** 2

(ii) the number of balls of color k assigned to a bin never exceeds the k-quota of the corresponding dimension.

The bins are chosen greedily: for each ball, the algorithm picks a bin with the maximum free space (the potential of the corresponding dimension minus the number of balls already put into it) among those which satisfy conditions (i) and (ii) above.

Theorem 5. *Algorithm 1 correctly calculates the geodetic Radon number of a given d-dimensional grid G in $\mathcal{O}(d\ (\log d)^{1/2})$ time.*

Proof. We prove the correctness of the proposed algorithm in two steps. First, we show that, if the algorithm fails to find a valid distribution of all

$$t(r) := \sum_{k=1}^{\lfloor r/2 \rfloor} template\_count_r(k)$$

balls into the d bins corresponding to the grid's dimensions, then no such distribution exists, in which case it is straightforward to conclude that there are no d ordered partitions of $[1, r]$ satisfying the requirements of Theorem 1, so G admits no anti-Radon set of size r. Second, we show that, if the algorithm does find such a distribution, then G admits an anti-Radon set of size r.

Let $\mathcal{D}$ be the non-empty collection of all *feasible distributions* of $t(r)$ colored balls into d bins, that is, distributions that satisfy the aforementioned conditions (i) and (ii). Let A denote the distribution obtained by the greedy algorithm, assuming, by contradiction, that $A \notin \mathcal{D}$. Distribution A is therefore incomplete (it was aborted by the algorithm midway through). We label the balls $1, 2, \ldots, b$ in the order they were considered by the algorithm, and we (re-)label the colors $1, 2, \ldots, c$ in the same fashion, according to the order in which they are considered by the algorithm. Let $A_m(j, k)$ denote the number of balls colored k placed in the j-th bin according to A among the first $m \leq b$ balls distributed by the greedy algorithm. Finally, for a feasible distribution $D \in \mathcal{D}$, let $D(j, k)$ be the number of balls with color k placed in the j-th bin according to D.

We define $m\colon \mathcal{D} \to [1, b]$ as the function

$$m(D) := \min \{m\colon A_m(j, k) > D(j, k) \text{ for some } j \in [1, d], k \in [1, c]\}.$$

To follow a standard "cut and paste" argument for proving the correctness of greedy algorithms, we let $D^* \in \mathcal{D}$ be chosen so that $m(D^*)$ is maximum. We call the ball whose label is $m(D^*)$ the *pivot* ball. Because D^* is a feasible distribution and, by hypothesis, A is not, there must exist a bin $j' \in [1, d]$ and a color $\hat{k} \in [1, c]$ such that the number of balls with color $\hat{k}$ put into bin j' according to A is greater than the number of balls colored $\hat{k}$ put into the same bin according to D, that is, $A_b(j', \hat{k}) > D(j', \hat{k})$. If that was not true, then A would not have been aborted in line 7 of the algorithm. Now, since $A(j', \hat{k}) > D^*(j', \hat{k})$ and the number of balls colored $\hat{k}$ in A is no greater than the number of such balls in D^*, there must be a bin $j'' \neq j'$ such that $A(j'', \hat{k}) < D^*(j'', \hat{k})$.

We now infer a contradiction by showing that it is possible to obtain a new feasible distribution D^{**} that is quite similar to D^* and satisfies $m(D^{**}) > m(D^*)$. We consider two cases:

(1) The bin j' is not full in D^*, that is, the number of balls put into bin j' by D^* is less than its potential.

(2) The bin j' is full in D^*.

If (1) holds, then it is simple to create a new feasible distribution D^{**} identical to D^* except that one ball with color $\hat{k}$ is moved from bin j'' to bin j'. We remark that, by doing so, the $\hat{k}$-quota of j' in D^{**} will not be exceeded, since the number of balls colored $\hat{k}$ in bin j' according to D^* was less than that same number according to A, and therefore strictly less than the maximum allowed.

If (2) holds, then we still want to move one ball with color $\hat{k}$ from bin j'' to bin j'. In this case, however, we must create room for it in bin j', so the number of balls does not exceed its potential. We now show that it is always possible to do so, because, in D^*, the number of balls with colors $k > \hat{k}$ in bin j' is strictly greater than the number of those balls in bin j''. This implies the existence of a color $\tilde{k}$ such that $D^*(j', \tilde{k}) \geq 1$ and $D^*(j'', \tilde{k})$ is less than the $\tilde{k}$-quota of j'', in which case one ball with color $\tilde{k}$ can be moved from bin j' to bin j'' creating the necessary space in bin j' for the ball colored $\hat{k}$ migrating from bin j''.

More formally, we want to show that

$$\sum_{k=\hat{k}+1}^{c} D^*(j', k) > \sum_{k=\hat{k}+1}^{c} D^*(j'', k).$$

Since bin j' is full in distribution D^*, we have

$$potential_r(j') = \sum_{k=1}^{\hat{k}-1} D^*(j', k) + D^*(j', \hat{k}) + \sum_{k=\hat{k}+1}^{c} D^*(j', k).$$

The j''-th bin, however, may not be full, and therefore we write

$$potential_r(j'') \geq \sum_{k=1}^{\hat{k}-1} D^*(j'', k) + D^*(j'', \hat{k}) + \sum_{k=\hat{k}+1}^{c} D^*(j'', k).$$

For $j \in [1, d]$, let $f(j)$ denote the free space in the j-th bin, according to A, by the time the greedy algorithm considers the placement of the pivot ball. It is easy to see that

$$f(j') = potential_r(j') - \sum_{k=1}^{\hat{k}-1} D^*(j', k) - D^*(j', \hat{k}),$$

since the number of balls in bin j', immediately before the placement of the pivot ball, is precisely the number of balls colored $k \leq \hat{k}$ in bin j' according to the feasible distribution D^* (distribution A performed by the algorithm had not deviated from D^* until that moment).

Similarly,

$$f(j'') \geq potential_r(j'') - \sum_{k=1}^{\hat{k}-1} D^*(j'', k) - A(j'', \hat{k}),$$

and the inequality is due to the fact that the final number $A(j'', \hat{k})$ of balls colored $\hat{k}$ in bin j'' according to A is an upper bound for the number of balls that would have already been placed by the algorithm in the j''-th bin before the pivot ball was considered.

From the fact that the algorithm chose the j'-th bin, not the j''-th (which certainly could get one more ball with color $\hat{k}$ at that point, as happens in D^*), we know that $f(j') \geq f(j'')$, implying

$$\sum_{k=\hat{k}+1}^{c} D^*(j', k) \geq D^*(j'', \hat{k}) + \sum_{k=\hat{k}+1}^{c} D^*(j'', k) - A(j'', \hat{k}).$$

Since $D^*(j'', \hat{k}) > A(j'', \hat{k})$, the desired result follows.

Now for the sufficiency. If the algorithm finds a feasible distribution of balls into bins as described, then there are d ordered partitions of $[1, r]$ as required by Theorem 1. The idea is that an ordered partition may eliminate templates $\langle R_1 \rangle, \langle R_2 \rangle, \ldots, \langle R_q \rangle$, with $R_1 \subsetneq R_2 \subsetneq \cdots \subsetneq R_q$, by letting R_1 be the first partite set, $R_2 \setminus R_1$ the second partite set, $R_3 \setminus R_2$ the third partite set, and so on, as illustrated in [1, proof of Lemma 5]. Moreover, the templates to be assigned to each ordered partition can be chosen by following the (almost-)perfect matchings admitted by the bipartite graphs with vertex sets $V_1 = \binom{[1,r]}{k}$ and $V_2 = \binom{[1,r]}{k+1}$ where $u \in V_1$ is adjacent to $v \in V_2$ if and only if $u \subseteq v$, as in Eckhoff's original proof for the Radon number of the Manhattan convexity on $\mathbb{R}^d$ [2].

As for the time complexity, each iteration of the outer loop looks for anti-Radon sets of size r, and its running time is clearly linear in the number of Radon partition templates that must be eliminated. Since the number of such partitions for a set of size r is $2^{r-1} - 1 = \mathcal{O}(2^r)$, the overall time complexity of the algorithm is

$$\sum_{r=2}^{r(d)-1} \mathcal{O}(2^r) = \mathcal{O}(2^{r(d)}).$$

We can now employ the approximation

$$\binom{r}{\lfloor \frac{r}{2} \rfloor} \approx \frac{2^r}{\sqrt{r+1}} \cdot \sqrt{\frac{2}{\pi}} \tag{3}$$

for the central binomial coefficient [5] to derive $r(d) = \mathcal{O}(\log d)$, which suffices to show that the time complexity of the algorithm is bounded from above by a polynomial in d. To make it more precise, we gather from (2) that

$$\binom{r(d)-1}{\left\lfloor \frac{r(d)-1}{2} \right\rfloor} \leq 2d, \tag{4}$$

and we replace r with $r(d) - 1$ in (3) before we plug it into (4), obtaining

$$\frac{2^{r(d)-1}}{\sqrt{r(d)}} \leq d\sqrt{2\pi} + c,$$

for some positive constant c. We can now write $2^{r(d)} r(d)^{-1/2} = \mathcal{O}(d)$, and the overall $2^{r(d)} = \mathcal{O}(d\ (\log d)^{1/2})$ time complexity follows. □

3.1. The Cartesian product of equal-sized paths

A variation of the proposed algorithm can be devised for the situation where all d dimensions of the grid G have the same size n. First of all, if $n = 2$, then the Radon number of G is given directly by $2 + \lfloor \log_2(d+1) \rfloor$; and if $n = 3$, then the Radon number is $3 + \lfloor \log_2 d \rfloor$. These results appeared in [1]. If $n \geq 4$, then, to find out whether the input grid admits an anti-Radon set of size r, the balls-into-bins phase of Algorithm 1 will simply put the first ball into the first bin, the second ball into the second bin, and so on up to the d-th bin. Then the $(d + 1)$-th ball will be put into the first bin again, the $(d + 2)$-th ball into the second been etc., provided the potential (as per Observation 2) and the k-quotas (as per Observation 4) of each dimension are respected. But now a simple calculation shall answer whether this process would succeed in distributing all balls, and *it does not have to be run at all.*

With respect to the potentials, the total number of balls $2^{r-1} - 1$ must not exceed $d \cdot (\min\{n, r\} - 1)$; and, with respect to the k-quotas, the number $\binom{r}{\lfloor r/2 \rfloor}$ of Radon partition templates of size $\lfloor r/2 \rfloor$ must not exceed $2d$. It is a simple exercise to check that these two conditions are also sufficient. The time complexity of the algorithm is thus $\mathcal{O}(r) = \mathcal{O}(\log d)$.[1]

4. Final considerations

The concluding section of [1] poses the question of whether the simple structure of d-dimensional grids can be exploited to allow for the efficient determination of its Radon number. This paper shows it can be done not only in polynomial time, but in near-linear time for arbitrary grids, and even in sub-linear time for those special grids that are the Cartesian products of paths of equal size. After implementing the proposed algorithm, we were able to replicate the complete results up to the ninth dimension given in [1] in a matter of milliseconds.[2]

A question that remains open is whether the Radon number can also be determined in polynomial time for (induced) subgraphs of grids, not only in the geodetic convexity but in other types of graph convexities as well.

References

[1] Dourado, M.C., Pereira de Sá, V.G., Rautenbach, D., Szwarcfiter, J.L., On the geodetic Radon number of grids, *Discrete Mathematics* **131** (2013), 111–121.

[2] Eckhoff, J., Der Satz von Radon in konvexen Produktstrukturen. II, *Monatsh. Math.* **73** (1969), 7–30.

[3] Gruber, P.M., Aspects of convexity and its applications, *Exposition. Math.* **2** (1984), 47–83.

[1] One may argue that, in this case, the input might be given as two numbers, d and n, with $\mathcal{O}(\log d)$ and $\mathcal{O}(\log n)$ bits, respectively. While sub-linear on the numerical value of the input, the complexity of the algorithm, in this case, is of course linear on its size.

[2] The Python code for general grids and for P_k^d can be found respectively in `https://www.dropbox.com/s/8z3xcqg6m63gi3j/linear_geodetic_radon_grids.py`, `https://www.dropbox.com/s/62cn4433qy00g1r/sublinear_geodetic_radon_Pkd.py`.

[4] Jamison-Waldner, R.E., Partition numbers for trees and ordered sets, *Pac. J. Math.* **96** (1981), 115–140.

[5] Koshy, T., *Catalan numbers with applications*, Oxford University Press (2008).

[6] Pelayo, I.M., *Geodesic Convexity in Graphs*, Springer (2013).

[7] Radon, J., Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten, *Mathematische Annalen* **83**:1–2 (1921), 113–115.

[8] Rosenfeld, M., Zaks, J. (eds.), *Convexity and Graph Theory*, Elsevier (1984).