Covering radii of 3-zonotopes and the shifted lonely runner conjecture

David Alcántara* Francisco Criado† Francisco Santos‡

Abstract

We show that the shifted Lonely Runner Conjecture (sLRC) holds for 5 runners. We also determine that there are exactly 3 primitive tight instances of the conjecture, only two of which are tight for the non-shifted conjecture (LRC). Our proof is computational, relying on a rephrasing of the sLRC in terms of covering radii of certain zonotopes (Henze and Malikiosis, 2017), and on an upper bound for the (integer) velocities to be checked (Malikiosis, Santos and Schymura, 2024+).

As a tool for the proof, we devise an algorithm for bounding the covering radius of rational lattice polytopes, based on constructing dyadic fundamental domains.

This is an extended abstract for the preprint [1].

1 Introduction

The lonely runner conjecture (LRC) states that if $n+1$ runners run along a circle of length one with constant, distinct, velocities, all starting at the origin, then for every runner there is a time at which all other runners are at distance at least $1/(n+1)$ from it. It was posed in 1968 by J. Wills [14] in the language of diophantine approximation, and is currently proved up to $n = 6$ [2]. The conjecture has attracted quite some attention due to the simplicity of its statement and because it admits various interpretations, from its original diophantine approximation statement, to visibility obstruction, billiard trajectories or nowhere zero flows in graphs, among others. See [12] for a very recent survey. We are interested in the so-called *shifted* version, a generalization in which runners are allowed to have different starting points. This version appeared in print for the first time in 2019 [3].

In both the original and the shifted versions, the runner we are looking at can be fixed at the origin, since only relative velocities are important. Hence the shifted conjecture becomes the following (the original LRC is the special case where $s_i = 0$ for all i):

Conjecture 1 (sLRC) *Let $v_1, \dots, v_n,\ s_1, \dots, s_n \in \mathbb{R}$ be real numbers, with the v_i distinct and non-zero. Then, there is a $t \in \mathbb{R}$ such that for every $i \in [n] := \mathbb{Z}/n\mathbb{Z}$, the distance between $v_i t + s_i$ and the closest integer, $\operatorname{dist}(v_i t + s_i, \mathbb{Z}) \geq \frac{1}{n+1}$.*

Since the order of the runners is not relevant, we assume without loss of generality that $v_1 < \cdots < v_n$.

This shifted version of the Lonely Runner Conjecture is only currently proved up to $n = 3$ (“four runners”) [5, 13].[1] For our proof we use that in Conjecture 1 (and in the original LR conjecture) there is no loss of generality in assuming all velocities v_i to be positive integers [5, 6, Section 4.1]. We then rely on the following result of Malikiosis, Santos and Schymura:

Theorem 1 ([11, Corollary 1.11]) *sLRC holds for $n = 4$ for all integer velocities with sum at least 196.*

That is, only the velocity vectors $(v_1, v_2, v_3, v_4) \in \mathbb{Z}$ with $1 \leq v_1 < v_2 < v_3 < v_4$ and $v_1 + v_2 + v_3 + v_4 \leq 195$ need to be checked. We can also assume $\gcd(v_1, v_2, v_3, v_4) = 1$ since dividing all velocities by a common factor c does not change the problem: the positions at time t of the original problem coincide with the positions at time ct of the new one. With these considerations our main result is:

Theorem 2 *There are 2 133 561 velocity vectors $(v_1, v_2, v_3, v_4) \in \mathbb{Z}$ with $1 \leq v_1 < v_2 < v_3 < v_4$, $v_1 + v_2 + v_3 + v_4 \leq 195$ and $\gcd(v_1, v_2, v_3, v_4) = 1$. The sLRC holds for all of them.*

Corollary 3 *sLRC (Conjecture 1) holds for $n = 4$ (five runners).*

We also show there are exactly 3 primitive integer velocity vectors that are *tight*, meaning that for them the bound $\frac{1}{5}$ is the best possible.

Theorem 4 *The only integer velocity vectors $(v_1, v_2, v_3, v_4) \in \mathbb{Z}^4$ with $1 \leq v_1 < v_2 < v_3 < v_4$, and $\gcd(v_1, v_2, v_3, v_4) = 1$ for which there are starting points $s_1, \dots, s_n \in \mathbb{R}$ such that for every time $t \in \mathbb{R}$, there is an index $i \in [n]$ such that $\operatorname{dist}(v_i t + s_i, \mathbb{Z}) \leq \frac{1}{5}$, are $(1, 2, 3, 4)$, $(1, 3, 4, 6)$, and $(1, 3, 4, 7)$.*

*Departmento de Matemáticas, Estadística y Computación, Universidad de Cantabria, Santander, Spain, david.alcantara@unican.es. Funded by grant PRE2022-000020 of MCIN/AEI/10.13039/501100011033.

†Departamento de Matemáticas, CUNEF Universidad, Madrid, Spain, francisco.criado@cunef.edu

‡Departmento de Matemáticas, Estadística y Computación, Universidad de Cantabria, Santander, Spain, francisco.santos@unican.es

[1] Observe that Rifford [13] refers to this as the case $n = 4$

Clearly, being tight for the nonshifted version implies tight for the shifted one. As an example, it is easy to show, as was already observed by Wills [15], that the vector $(1,\dots,n)$ is tight for the nonshifted version, for all $n \in \mathbb{N}$: suppose that at a certain time t we have no $v_i t$ in $[-1/(n+1), 1/(n+1)] + \mathbb{Z}$. Then the pigeon-hole principle implies that at that time at least two runners, say i and i' must be in the same interval $[j/(t+1),(j+1)/(t+1)] + \mathbb{Z}$ for some $j \in [n-1]$. But then, assuming w.l.o.g. that $i' > i$, we have the contradiction that[2]

$$v_i t, v_{i'} t \in \left[\tfrac{j}{n+1}, \tfrac{j+1}{n+1}\right] + \mathbb{Z} \qquad \Rightarrow$$
$$v_{i'-i} t = v_{i'} t - v_i t \in \left[-\tfrac{1}{n+1}, \tfrac{1}{n+1}\right] + \mathbb{Z},$$

The vectors $(1,2,3,4)$ and $(1,3,4,7)$ are known to be tight for the non-shifted LRC and, in fact, they are the only ones for $n=4$ [7]. The vector $(1,3,4,6)$ is new and shows that shifted tightness does not imply the unshifted one.

Our method (as well as the proof of Theorem 1 in [11]) is based on the relation between the Lonely Runner conjecture (both shifted and original one) to $(n-1)$-dimensional zonotopes with n generators [3, 8]. In particular, the sLRC can be restated as a bound on the covering radius of a certain class of zonotopes in $\mathbb{R}^{n-1}$.

Our proofs of Theorems 2 and 4 are computational; for each primitive velocity vector we build a fundamental domain of the integer lattice $\mathbb{Z}^3$ that fits in the dilated zonotope associated to it. This certifies that its covering radius satisfies the bound.

To prove tightness of the instances of Theorem 4, we explicitly find the *last covered points* of the zonotopes. These points correspond to the starting points of the sLRC which are tight for those velocity vectors.

Our algorithm to construct fundamental domains can in fact decide the covering radius of arbitrary lattice polytopes in any dimension.

2 Zonotopal statement of the LRC

We here recall the reformulation of Conjecture 1 in terms of zonotopal geometry, derived in [3, 8, 11].

A *zonotope* is any Minkowski sum of line segments. As such, any zonotope Z can be written as

$$\mathbf{c} + \sum_{i=1}^{n} [\mathbf{0}, \mathbf{u}_i] = \left\{ \mathbf{c} + \sum_{i=1}^{n} \lambda_i \mathbf{u}_i : \lambda_i \in [0,1]\ \forall i \right\},$$

for a certain finite set $\mathbf{u}_1,\dots,\mathbf{u}_n \in \mathbb{R}^d$ of vectors, called the *generators* of Z, and a certain point $\mathbf{c}$. This point is not important for us, since all that we do is invariant under translation. One natural choice is $\mathbf{c} = \mathbf{0}$ but often a more convenient choice is $\mathbf{c} = \frac{1}{2}\sum_{i=1}^{n} \mathbf{u}_i$, since it makes the zonotope become $Z = \frac{1}{2}\sum_{i=1}^{n}[-\mathbf{u}_i, \mathbf{u}_i]$, and be centrally symmetric around the origin.

[2] This argument is an instance of the general proof of Dirichlet's approximation theorem.

2.1 Lonely runner zonotopes and volume vectors

Definition 1 *A* Lonely Runner (LR) Zonotope *is any zonotope $Z \subset \mathbb{R}^{n-1}$ generated by a set of n integer vectors $\mathbf{U} = \{\mathbf{u}_i : 1 \le i \le n\} \subset \mathbb{Z}^{n-1}$ in linear general position; that is, such that every $n-1$ of them are a linear basis of $\mathbb{R}^{n-1}$.*

The volume vector *of Z is the vector $\mathbf{v} = (v_1,\dots,v_n) \in \mathbb{Z}_{>0}$ defined by*

$$v_i := |\det(\mathbf{U} \setminus \{\mathbf{u}_i\})|. \tag{1}$$

When all entries of the volume vector are distinct we say that Z is a strong Lonely Runner (sLR) Zonotope.

We call $\mathbf{v}$ the *volume vector* of Z, because its entries are the volumes of the n parallelepipeds that make up Z. In particular we have that $\mathrm{vol}(Z) = \sum_{i=1}^{n} v_i$ (see details, e.g., in [11]). Observe also that the generators and the volume vector satisfy

$$v_1\mathbf{u}_1 \pm \dots \pm v_n \mathbf{u}_n = 0$$

for some choice of signs. In fact, this equation (together with positivity of the v_i) characterizes $\mathbf{v}$ for given generators, modulo a scalar factor.

In the following result and the rest of the paper, a unimodular transformation is an affine transformation with integer coefficients and determinant ± 1.

Proposition 5 (Prop. 2.2, [1], see also §1.2, [11]) *For every integer $\mathbf{v} = (v_1,\dots,v_n) \in {\mathbb{Z}_{>0}}^n$ there is some LR zonotope with integer generators and with volume vector $\mathbf{v}$. If $\gcd(v_1,\dots,v_n) = 1$, then any two such zonotopes are equivalent by a unimodular transformation.*

2.2 Covering radius and the sLRC

A *convex body* in $\mathbb{R}^d$ is a convex compact subset. We assume our convex bodies to be *nondegenerate*, that is, that they have non-empty interior or, equivalently, that they are not contained in a hyperplane. This includes all bounded full-dimensional polytopes.

Definition 2 (Covering radius) *Let $C \subseteq \mathbb{R}^d$ be a convex body. The* covering radius of C, *denoted $\mu(C)$, is the smallest dilation factor $\rho > 0$ such that*

$$\rho C + \mathbb{Z}^d = \mathbb{R}^d.$$

The covering radius is invariant under real translations and unimodular transformations of C since they amount to translations and unimodular transformations of $\rho C + \mathbb{Z}^d$.

The zonotopal restatement of the *Shifted Lonely Runner Conjecture* is the following. Our statement is taken from [11] but the result is implicit in [3,5,8].

Proposition 6 ([8], see also [11, Proposition 1.8]) *Let $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{Z}^n_{>0}$ with pairwise distinct entries and $\gcd(v_1, \ldots, v_n) = 1$. Then, the following are equivalent:*

1. *A time t as required by the* Shifted Lonely Runner Conjecture *exists for velocities $\mathbf{v}$.*
2. *The sLR zonotope Z with volume vector v has $\mu(Z) \le \frac{n-1}{n+1}$.*

2.3 Covering radius via fundamental domains

In order to apply this result one does not need to compute $\mu(Z)$ (which is quite expensive, see e.g. [5]), but only check whether a certain number is a bound for it. This checking is closely related to finding a fundamental domain inside a scaled copy of Z.

Recall that a *fundamental domain* of $\mathbb{R}^d$ (with respect to $\mathbb{Z}^d$) is a set containing exactly one representative of each coset $\mathbf{p} + \mathbb{Z}^d$, $\mathbf{p} \in \mathbb{R}^d$.The definition of covering radius trivially translates to:

Lemma 7 *Let C be a convex body in $\mathbb{R}^d$ and $\rho > 0$. Then, the following are equivalent:*

1. *$\mu(C) \le \rho$.*
2. *ρC contains representatives of all $\mathbb{R}^d/\mathbb{Z}^d$.*
3. *ρC contains a fundamental domain.*
4. *No open set $W \in \mathbb{R}^d$ satisfies $\rho C \cap (W + \mathbb{Z}^d) = \emptyset$.*

Proof. The implications (1) $\Leftrightarrow$ (2) $\Leftrightarrow$ (3) and (2) $\Rightarrow$ (4) are obvious. Let us prove (4) $\Rightarrow$ (1). $\mu(C) > \rho$ implies there is a $\mathbf{p} \notin \rho C + \mathbb{Z}^d$. Since $\rho C + \mathbb{Z}^d$ is closed, its complement $W := \mathbb{R}^d \setminus \rho C$ is open and

$$\rho C \cap (W + \mathbb{Z}^d) \subset (\rho C \cap (\mathbb{R}^d \setminus \rho C)) + \mathbb{Z}^d = \emptyset. \qquad \square$$

2.4 The denominator of the covering radius

It is well-known and easy to show that the covering radius of a rational polytope is rational (see, e.g., [10, Proposition 5.1]). We give an explicit bound for its denominator in terms of the defining equations. The denominator of a rational number ρ is defined as the minimum positive integer s such that $s\rho$ is an integer. Having a bound for the denominator allows our algorithms to certify an exact upper bound for $\mu(P)$ from an approximate one, as follows.

Proposition 8 *Let P be a rational polytope and let $D \in \mathbb{N}$ be an upper bound for the denominator of $\mu(P)$. (For example, but not necessarily, a bound obtained by Corollary 10). Let $\rho = r/s$ with $r, s \in \mathbb{Z}$ and $s > 0$. Then, the following equivalences hold:*

1. *$\mu(P) \le \rho$ if and only if $\mu(P) < \rho + \frac{1}{sD}$*
2. *$\mu(P) \ge \rho$ if and only if $\mu(P) > \rho - \frac{1}{sD}$*

Proof. One direction is obvious in both cases. For the other one, we know that $\mu(P) = \frac{r'}{s'}$ for integers r', s' with $0 < s' \le D$. Assuming $\frac{r'}{s'} \ne \frac{r}{s}$ we have that

$$|\mu(P) - \rho| = \left|\frac{r'}{s'} - \frac{r}{s}\right| = \left|\frac{r's - rs'}{s's}\right| \ge \frac{1}{ss'} \ge \frac{1}{sD}.$$

Hence, either $\mu(P) = \rho$, $\mu(P) \ge \rho + \frac{1}{sD}$ or $\mu(P) \le \rho - \frac{1}{sD}$. $\square$

Our bound uses the concept of *last covered point*.

Definition 3 (Last covered point [4,5]) *Let $C \subseteq \mathbb{R}^d$ be a convex body. A* last covered point *for C is any $\mathbf{p} \in \mathbb{R}^d$ with $\mathbf{p} \notin (\mu(C)C)^\circ + \mathbb{Z}^d$.*

Since $\mu(C)$ is invariant under translation, we may assume $\mathbf{0} \in C^\circ$ without loss of generality. This simplifies some arguments because it implies that $\rho C + \mathbf{q}$ is monotone increasing in ρ. Therefore, once a point is covered by some copy $\rho C + \mathbf{q}$, it is also covered by the same copy for any larger ρ'.

In particular, under the assumption $\mathbf{0} \in C^\circ$ we have that a point $\mathbf{p}$ is last covered if $\rho C + \mathbb{Z}^d$ does not contain $\mathbf{p}$ for any $\rho < \mu(C)$, which explains the name.

Observation 1 *The set of last covered points is always non-empty.*

Proof. Assuming without loss of generality $\mathbf{0} \in C^\circ$, for each point $\mathbf{p} \in \mathbb{R}^d$ let $\rho_{\mathbf{p}} = \min\{\rho \in \mathbb{R}_{\ge 0} : \mathbf{p} \in (\rho C)^\circ + \mathbb{Z}^d\}$ be the *covering time* of the point $\mathbf{p}$. This definition is invariant by integer translations, and continuous. Since $\mathbb{R}^d/\mathbb{Z}^d$ is compact, there must be some point $\mathbf{p}$ with maximal covering time, i.e. $\rho_{\mathbf{p}} = \mu(C)$. This point is a last covered point. $\square$

In the rest of the section, $P \in \mathbb{R}^d$ is a polytope defined by the system of inequalities $Ax \le b$ for some matrix $A \in \mathbb{R}^{m\times d}$ and vector $b \in \mathbb{R}^m$. For an element $i \in [m]$ or subset $I \subset [m]$, A_i, b_i, A_I, b_I, etc. denote the restriction of a matrix or vector to the rows labelled by i or I.

Lemma 9 ([5, Lemma 3.1]) *Let $P = \{Ax \le b\}$ be a polytope and let $\rho = \mu(P)$. Then, there is*

- *a subset $R \subset [m]$ of rows with $|R| = d + 1$ and $\det(A_R|b_R) \ne 0$ and*
- *a lattice point $\mathbf{q}_i \in \mathbb{Z}^d$ for each $i \in R$,*

such that the system

$$A_i(\mathbf{x} - \mathbf{q}_i) = \rho b_i \quad \forall i \in R \tag{2}$$

has a unique solution in $\mathbb{R}^{d+1}$ and this solution is a last covered point.

The proof uses an extremal argument and Farkas' lemma. See Section A.1 in the appendix for details.

Corollary 10 *Let P be a rational polytope described by $A\mathbf{x} \leq b$ with $A \in \mathbb{Z}^{m\times d}$ and $b \in \mathbb{Z}^m$. Then $\mu(P)$ is a rational number and its denominator is bounded by*

$$\max_{R\in\binom{[m]}{d+1}} |\det(A_R|b_R)|.$$

Proof. Apply Cramer's rule to the variable t in the system of Lemma 9. □

3 Our Algorithms

In this section we introduce the algorithm we have used to find fundamental domains within each relevant sLR zonotope, postponing some details until Section 4.

Obtaining a representative zonotope for a velocity vector is discussed in [11]. We further simplify our representatives, reducing the length of their generators using the LLL algorithm, as described in Section 4.1.

3.1 Certifying an upper bound for the cov. radius

We here describe an algorithm to decide whether a facet-defined polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \leq b$ contains a fundamental domain. By Lemma 7, this is equivalent to certifying a given upper bound ρ for the covering radius of a polytope.

We consider a special family of fundamental domains of the integer lattice, given by unions of *dyadic voxels.*

Definition 4 *A* dyadic d-voxel *of* level $\ell \in \mathbb{Z}_{\geq 0}$ *is a half-open cube of the form*

$$\mathbf{c} + \frac{1}{2^\ell}[0,1)^d,$$

for some dyadic point $\mathbf{c} \in \frac{1}{2^\ell}\mathbb{Z}^d$. The integer point $\lfloor \mathbf{c} \rfloor$ is the displacement *of the voxel, and the difference $2^\ell(\mathbf{c} - \lfloor \mathbf{c} \rfloor)$ is the* type *of the voxel.*

All dyadic voxels of the same type are equivalent by integer translation and the voxel types are naturally arranged as an infinite rooted 2^ℓ-ary tree with the voxels of level ℓ at depth ℓ. We call this the infinite dyadic tree.[3]

A *dyadic fundamental domain* is a fundamental domain obtained as a finite union of dyadic voxels.

Every dyadic fundamental domain can be expressed as (the leaves of) a full-subtree of the infinite dyadic tree, with leaves labelled by their displacements.

[3]One can represent each type of level ℓ as a vector $\mathbf{b} = (\mathbf{b}_1, \ldots, \mathbf{b}_d)$ where each $\mathbf{b}_i$ is a binary string of length ℓ. In this representation $\mathbf{b}$ is an ancestor of $\mathbf{b}'$ in the infinite dyadic tree if and only if each $\mathbf{b}_i$ is an initial segment, or prefix, of the corresponding $\mathbf{b}'_i$. Equivalently, if the voxel of type $\mathbf{b}$ with zero displacement is contained in that of type $\mathbf{b}'$.

The simplest of such domains is the unit cube, that is, the root of the dyadic tree. Our algorithm performs a search in the infinite dyadic tree, starting with the root and iteratively subdividing all leaves which cannot be translated to fit in our zonotope, until either (a) all leaves fit inside, which certifies that we have constructed a dyadic fundamental domain contained in P or (b) the center of one leaf is found to have no translation inside P, certifying that no such fundamental domain exists.

This algorithm is illustrated in Figure 1.

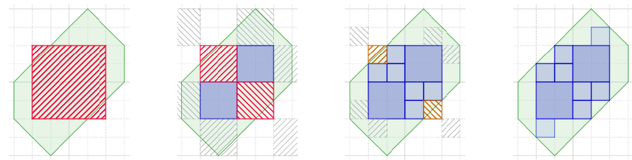

Figure 1: States of our Algorithm at different depths, applied to $\frac{1}{2}Z$ where Z is the 2-dimensional sLR zonotope with volume vector $(1, 2, 4)$. Notice that there are two choices for each voxel in the last step.

The decision of whether a voxel admits an integer translation that fits in our zonotope requires checking the feasibility of an integer linear program.

Proposition 11 *Let $P = \{A\mathbf{x} \leq \mathbf{b}\} \subset \mathbb{R}^d$ be a polytope and let $V = \mathbf{c} + [0,\epsilon)^d$ be a voxel. Then, P contains an integer translation of V if and only if the following Integer Linear Program is feasible:*

$$\text{find } \mathbf{x} \in \mathbb{Z}^d \quad \text{subject to } A\mathbf{x} \leq \mathbf{b} - A\mathbf{c} - A_{\geq 0}\,\boldsymbol{\epsilon}, \quad (3)$$

where $\boldsymbol{\epsilon} \in \mathbb{R}^d$ is the vector with all entries equal to ϵ and $A_{\geq 0}$ denotes the matrix with entry (i,j)-th entry equal to $\max\{0, A_{ij}\}$, for every (i,j).

Proof. A half-empty voxel is contained in P if and only the closed voxel is, that is, if and only of all the vertices of the translated closed voxel are in P. Hence, for a given $\mathbf{x} \in \mathbb{Z}^d$, the voxel $\mathbf{x} + V = \mathbf{x} + \mathbf{c} + [0,\epsilon)^d$ is contained in P if and only if all the points $\mathbf{y} \in \{0,\epsilon\}^d$ satisfy the inequalities $A(\mathbf{y} + \mathbf{c} + \mathbf{x}) \leq \mathbf{b}$. Now, for each row A_i of A, the maximum value of the functional A_i on the set $\{0,\epsilon\}^d$ is precisely $(A_i)_{\geq 0}\,\boldsymbol{\epsilon}$. □

The search algorithm as described so far has one issue: P may contain a fundamental domain but no dyadic one, and in this case our algorithm does not terminate. For rational polytopes we can solve this issue thanks to Proposition 8.

Theorem 12 *Let P be a rational polytope and let D be an upper bound for the denominator of $\mu(P)$. Let $\rho = r/s$ with $r, s \in \mathbb{Z}$ and $s > 0$.*

1. *If $\mu(P) \leq \rho$ then $\left(\rho + \frac{1}{2sD}\right)P$ contains a dyadic fundamental domain.*

2. *If $\mu(P) > \rho$ then there is an $\ell \in \mathbb{Z}_{\geq 0}$ and a dyadic point $\mathbf{c} \in \frac{1}{2^\ell}\{0, \dots, 2^\ell - 1\}^d$ such that $\left(\rho + \frac{1}{2sD}\right) P$ does not intersect $\mathbf{c} + \mathbb{Z}^d$.*

See Section 4.2 for a proof.

Hence, rather than applying our algorithm to the zonotopes dilated by $\frac{3}{5}$ we dilate them by $(\frac{3}{5} + \frac{1}{10D})$, which does not affect the result, yet ensures the algorithm terminates.

4 Algorithm details

4.1 Enumeration, construction, and preprocessing of sLR zonotopes

According to Theorem 1 we only need to enumerate sLR zonotopes up to volume 195. We first construct the list of possible volume vectors, that is, the 4-tuples $v = (v_1, \dots, v_4) \in \mathbb{Z}^4$ with $0 < v_1 < v_2 < v_3 < v_4$. As observed in the introduction we can assume that $\gcd(v_1, v_2, v_3, v_4) = 1$. Moreover, by Proposition 5, with this restriction there is a unique sLR zonotope (modulo unimodular equivalence) for each volume vector. Enumerating such 4-tuples is algorithmically trivial and took less than a second in a standard PC:

Proposition 13 *There are exactly* 2 133 561 *vectors* $(v_1, v_2, v_3, v_4) \in \mathbb{Z}$ *with* $1 \leq v_1 < v_2 < v_3 < v_4$, $\gcd(v_1, v_2, v_3, v_4) = 1$ *and* $\sum v_i \leq 195$.

We then need to generate a representative zonotope from its volume vector v. This is done with Algorithm 1, which follows the 'existence' part of the proof of Proposition 5 given in [11].

Algorithm 1: Compute generators for a LR zonotope from its volume vector.

Input : $v = (v_1, \dots, v_n) \in \mathbb{Z}^n_{>0}$, with $\gcd(v_1, \dots, v_n) = 1$.

Output: A matrix $M = (\mathbf{u}_1, \dots, \mathbf{u}_n) \in \mathbb{Z}^{(n-1)\times n}$ such that $\mathbf{u}_1, \dots, \mathbf{u}_n$ generate a LR zonotope with volume vector v.

1 Let $M' := \left(\begin{array}{ccc|c} -v_n & & & v_1 \\ & \ddots & & \vdots \\ & & -v_n & v_{n-1} \end{array}\right)$.

2 Let $H \in \mathbb{Z}^{(n-1)\times n}$ be the column-wise Hermite normal form of M, and let $B \in \mathbb{Z}^{(n-1)\times(n-1)}$ consist of the first $n-1$ columns of H.

3 Apply an LLL-reduction to the rows of $B^{-1}M'$ and let $M \in \mathbb{Z}^{(n-1)\times n}$ have as rows the resulting reduced vectors.

4 **return** M.

Step 1 in the algorithm creates an integer matrix $M' \in \mathbb{Z}^{(n-1)\times n}$ whose columns generate a LR zonotope with volume vector a scalar multiple of $(v_1, \dots, v_n)$. Step 2 then uses a column-wise Hermite normal form of M' to construct a basis (the columns of the matrix B in the algorithm) of the lattice Λ generated by the $\mathbf{u}'_i$. Observe that $\mathrm{rk}(M') = n - 1$ implies that the last column of its Hermite normal form H is zero, and B is simply equal to H without that column.

Now, B^{-1} is the matrix of a linear isomorphism $\Lambda \xrightarrow{\cong} \mathbb{Z}^{n-1}$, so the columns of $B^{-1}M'$ would already be valid generators for a LR zonotope with volume vector $(v_1, \dots, v_n)$. However, the generators obtained in this way typically have some large entries, resulting in 'long and skinny' zonotopes that are are poorly conditioned for our method to compute covering radii. To overcome this we preprocess the generators in step 3 of the algorithm, by performing an LLL lattice basis reduction to the rows of $B^{-1}M'$.[4] This produces a matrix M whose columns are unimodularly equivalent to those of $B^{-1}M'$, but with smaller entries.

For our covering radius computations we need to convert the generators of the zonotope into an inequality description of it. This, for an arbitrary zonotope $Z \subset \mathbb{R}^d$ with generators $U = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ is done as follows, where we are identifying $\bigwedge^{d-1} \mathbb{R}^d \cong (\mathbb{R}^d)^*$ in the natural way.

Proposition 14 *Let* $Z = \frac{1}{2}\sum_{i=1}^n [-\mathbf{u}_i, \mathbf{u}_i]$ *be the* $\mathbf{0}$*-symmetric zonotope with generators* $\mathbf{u}_1, \dots, \mathbf{u}_n$*. Then*

$$Z = \left\{ \mathbf{x} \in \mathbb{R}^d : -b_S \leq a_S \mathbf{x} \leq b_S : S \in \binom{[n]}{d-1} \right\},$$

where

$$a_S := \bigwedge_{i \in S} \mathbf{u}_i \in (\mathbb{R}^d)^* \qquad \text{and} \qquad b_S := \frac{1}{2}\sum_{i=1}^n |a_S\, \mathbf{u}_i|.$$

Proof. Each facet of a zonotope is a zonotope itself, generated by the $\mathbf{u}_i$ contained in, and spanning, a linear hyperplane. Hence, every normal vector is indeed of the form a_S for some $(d-1)$-subset S of U.

By central symmetry, there are two parallel facets with normal vectors $\pm a_S$. The corresponding facet inequalities are $-b_S \leq a_S\mathbf{x} \leq b_S$, since $\pm b_S$ are the minimum and maximum values taken by a_S in the set

$$\left\{ \sum_{i=1}^n \pm \mathbf{u}_i \right\},$$

which contains all vertices of Z. □

4.2 Building a dyadic fundamental domain

In this section we present Algorithm 2, the concrete algorithm that explores the infinite dyadic tree to decide

[4] We have implemented the LLL algorithm with $\delta = 3/4$. Higher values of $\delta(0, 1)$ would give better zonotopes, but would increase the running time.

the covering radius of an arbitrary lattice polytope, as discussed in Section 3.

The algorithm requires a facet description of the polytope, which can be derived from the generators of a zonotope by Proposition 14.

The termination of this algorithm follows from Theorem 12, of which we give now proof.

Proof. (of Theorem 12) For part (1) we only need to use that

$$\mu(P^+) = \frac{\mu(P)}{\rho + \frac{1}{2sD}} \leq \frac{\rho}{\rho + \frac{1}{2sD}} < 1.$$

For each $\ell \in \mathbb{N}$ let D_ℓ be the union of all the dyadic boxes of depth ℓ contained in P^+. Since D_ℓ converges (e.g. in the Hausdorff metric) to P^+ when ℓ goes to infinity, we have that $\mu(D_\ell)$ converges to $\mu(P^+)$. In particular, there is an ℓ such that $\mu(D_\ell) < 1$. Hence, D_ℓ contains a fundamental domain, and this fundamental domain can be obtained taking one representative for each type of voxel in the union D_ℓ.

For part (2) we use that $\mu(P) > \rho$ implies (by Corollary 8) that $\mu(P) \geq \rho + \frac{1}{sD}$. Hence

$$\mu(P^+) = \frac{\mu(P)}{\rho + \frac{1}{2sD}} \geq \frac{\rho + \frac{1}{sD}}{\rho + \frac{1}{2sD}} > 1.$$

The statement then follows from the density of the dyadic points $\mathbb{Z}[\frac{1}{2}]^d$ in $\mathbb{R}^d$ and Lemma 7, which asserts the existence of an open set $W \subset \mathbb{R}^d \setminus (P^+ + \mathbb{Z}^d)$. □

Theorem 15 *Algorithm 2 always terminates and it correctly decides whether $\mu(P) \leq \rho$ for any lattice polytope P and $\rho \in \mathbb{Q}_+$.*

Proof. Observe that the algorithm returns a certificate in either case. Let us first show their correctness.

If the algorithm finishes with a set of dyadic voxels, these voxels are a full subtree of the infinite dyadic tree by construction, and hence they form a dyadic fundamental domain. Furthermore, all of these voxels are contained in P^+, so $\mu(P) \leq (\rho + \frac{1}{2sD})$ and Corollary 8 implies $\mu(P) \leq \rho$.

On the other hand, if the algorithm finishes with a point $\mathbf{c}$ such that P^+ does not intersect $\mathbf{c}+\mathbb{Z}^d$, Lemma 7 implies $\mu(P) > (\rho + \frac{1}{2sD}) > \rho$.

To prove that the algorithm terminates we handle the two cases separately.

If $\mu(P) \leq \rho$, Theorem 12 guarantees the existence of a dyadic fundamental domain D contained in P^+. Let ℓ be the maximum depth of the voxels in D. Then every voxel type of depth $\geq \ell$ has a representative contained in P^+, so the algorithm will never enter the "else" in line 14 with a voxel of depth $\geq \ell$. Hence, the algorithm can perform the while loop only finitely many times before N becomes empty.

Algorithm 2: Decide whether $\mu(P) \leq \rho$.

Input : A rational polytope $P = \{A\mathbf{x} \leq \mathbf{b}\}$ (with A and $\mathbf{b}$ integer) and a rational number $\rho = r/s$, with $r, s \in \mathbb{Z}_+$.

Output: A dyadic fundamental domain S or a dyadic point $\mathbf{c}$ certifying whether $\mu(P) \leq \rho$ or not, as in Theorem 12.

1 Let D be a bound on the denominator of $\mu(P)$, such as $\mu_{det}(A, \mathbf{b})$.

2 Let

$$P^+ = \left(\rho + \frac{1}{2sD}\right) P = \left\{A\mathbf{x} \leq \left(\rho + \frac{1}{2sD}\right)\mathbf{b}\right\}$$

3 Initialise a queue N of 'nodes to be processed' containing the unit cube

4 Initialise an empty list S of 'voxels in the fundamental domain'

5 **while** *there are nodes in N* **do**

6   Let $V = \mathbf{c} + [0, \frac{1}{\ell^d})^d$ be one such node of maximum size.

7   Delete V from N and

8   **if** *P^+ does not intersect* $\mathbf{c} + \mathbb{Z}^d$ **then**

9   **return c**

10   **else**

11   **if** $\exists\, \mathbf{p} \in \mathbb{Z}^d$ *with* $\mathbf{p} + V \subset P^+$ **then**

12   add the voxel $\mathbf{p} + V$ to S

13   **else**

14   add the 2^d children of V to N

15 **return** S

If $\mu(P) > \rho$, Theorem 12 guarantees the existence of a dyadic point $\mathbf{c}$ with $(\mathbf{c} + \mathbb{Z}^d) \cap P^+ = \emptyset$. Let ℓ be the minimal depth of such a point. Since the algorithm processes the infinite dyadic tree in a breadth-first search manner, in a finite number of steps it will check all the dyadic points of depth ℓ (either implicitly for those contained in voxels of depth $\leq \ell$ and with $\mathbf{p} + V \subset P^+$, or explicitly for those not contained in such voxels). □

References

[1] D. Alcántara, F. Criado, F. Santos. Covering radii of 3-zonotopes and the shifted Lonely Runner Conjecture, preprint, 2025. https://arxiv.org/abs/2506.13379

[2] J. Barajas and O. Serra. The lonely runner with seven runners. *Electron. J. Combin.*, 15:#48, 18 pp. (electronic), 2008.

[3] M. Beck, S. Hoşten, and M. Schymura. Lonely Runner Polyhedra. *Integers*, 19:#A29, 13 pp, 2019.

[4] G. Codenotti, F. Santos, and M. Schymura. The covering radius and a discrete surface area for non-hollow simplices. *Discrete & Computational Geometry*, 67(1):65–111, 2022.

[5] J. Cslovjecsek, R. D. Malikiosis, M. Naszódi, and M. Schymura. Computing the covering radius of a polytope with an application to lonely runners. *Combinatorica*, 42(4):463–490, Feb. 2022.

[6] T. W. Cusick. View-obstruction problems. *Aequationes Math.*, 9:165–170, 1973.

[7] T. W. Cusick and C. Pomerance. View-obstruction problems III. *J. Number Theory*, 19:131–139, 1984.

[8] M. Henze and R. D. Malikiosis. On the covering radius of lattice zonotopes and its relation to view-obstructions and the lonely runner conjecture. *Aequationes Math.*, 91(2):331–352, 2017.

[9] Q. Huangfu and J. A. J. Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142, Dec. 2017.

[10] R. Kannan. Lattice translates of a polytope and the frobenius problem. *Combinatorica*, 12(2):161–177, June 1992.

[11] R. D. Malikiosis, F. Santos, and M. Schymura. Linearly-exponential checking is enough for the lonely runner conjecture and some of its variants, preprint, 2024. https://arxiv.org/abs/2411.06903v2

[12] G. Perarnau and O. Serra. The Lonely Runner Conjecture turns 60, preprint, 2024. https://arxiv.org/abs/2409.20160.

[13] L. Rifford. On the time for a runner to get lonely, 2021. *Acta Appl. Math.*, 180, 15, 2022.

[14] J. M. Wills. Zur simultanen homogenen diophantischen Approximation. I. *Monatsh. Math.*, 72:254–263, 1968.

[15] J. M. Wills. Zur simultanen homogenen diophantischen Approximation. II. *Monatsh. Math.*, 72:368–381, 1968.

Appendix

A.1 Bound for the denominator of the cov. radius

Proof. (of Lemma 9) As in Observation 1, we may assume $\mathbf{0} \in P^\circ$ without loss of generality, or equivalently, that $b_i > 0$ for all $i \in [m]$.

Each facet of a translated polytope $\{\rho P + \mathbf{q} : \mathbf{q} \in \mathbb{Z}^d\}$ is labelled by a point $\mathbf{q} \in \mathbb{Z}^d$ and an index $i \in [m]$. For each last covered point $\mathbf{p}$, let $R_\mathbf{p}$ be the set of indices i such that $\mathbf{p}$ lies in the i-th facet of $\rho P + \mathbf{q}_i$, for some $\mathbf{q}_i \in \mathbb{Z}^d$. Observe that $\mathbf{p}$ lies in the affine subspace $L_\mathbf{p} := \{\mathbf{x} \in \mathbb{R}^d : \forall i \in R_\mathbf{p}, A_i(\mathbf{x} - \mathbf{q}_i) = \rho b_i\}$, since $A_i(\mathbf{x} - \mathbf{q}_i) = \rho b_i$ is the facet equation for the i-th facet of $\rho P + \mathbf{q}_i$.

Choose a last covered point $\mathbf{p}$ so that $R_\mathbf{p}$ is maximal. Maximality implies that $L_\mathbf{p} = \{\mathbf{p}\}$, since otherwise moving the point within $L_\mathbf{p}$ until an extra facet of some $\{\rho P + \mathbf{q} : \mathbf{q} \in \mathbb{Z}\}$ is met (which happens at the latest when we are about to leave a certain $\rho P + \mathbf{q}_i$ containing $\mathbf{p}$) gives us a last covered point $\mathbf{p}'$ with $R_{\mathbf{p}'}$ strictly containing $R_\mathbf{p}$.

The fact that $L_\mathbf{p} = \{\mathbf{p}\}$ implies that the matrix $A_{R_\mathbf{p}}$ consisting of rows used in $R_\mathbf{p}$ has full rank, equal to d.

Now, observe that the vectors A_i for $i \in R_\mathbf{p}$ must have a positive linear dependence. Otherwise, by (one of many versions of) Farkas' lemma, there is a vector $\mathbf{v} \in \mathbb{R}^d$ such that $\langle A_i, \mathbf{v} \rangle > 0$ for all $i \in R_\mathbf{p}$. Then, $\mathbf{p}$ would not be last covered, as $\mathbf{p} + \varepsilon\mathbf{v}$ would not be covered by any $P + \mathbf{q}_i$. Locally, these are the only translated copies of P that could potentially cover $\mathbf{p}$. Therefore, no translated copy of P covers $\mathbf{p} + \varepsilon\mathbf{v}$, so $\mathbf{p} + \varepsilon\mathbf{v}$ would have larger covering time than $\mathbf{p}$, contradicting the assumption of $\mathbf{p}$ being last-covered.

The positive linear dependence of the vectors A_i for $i \in R_\mathbf{p}$ implies that the system of equalities

$$A_i(\mathbf{x} - \mathbf{q}_i) = tb_i, \quad i \in R_\mathbf{p},$$

where t is considered an extra variable, has no solution with $t \in [0, \rho)$. Indeed, let $\lambda_i \in \mathbb{R}_{\geq 0}$ for $i \in R_\mathbf{p}$ be the coefficients of the linear dependence we defined above. Then,

$$\begin{aligned} A_i((\mathbf{x} - \mathbf{p}) + \mathbf{p} - \mathbf{q}_i) &= (t - \rho + \rho)b_i, && i \in R_\mathbf{p}, \\ A_i(\mathbf{x} - \mathbf{p}) + A_i(\mathbf{p} - \mathbf{q}_i) &= \rho b_i + (t - \rho)b_i, && i \in R_\mathbf{p}, \\ A_i(\mathbf{x} - \mathbf{p}) &= (t - \rho)b_i, && i \in R_\mathbf{p}, \end{aligned}$$

$$\begin{aligned} \sum_{i \in R_\mathbf{p}} \lambda_i A_i(\mathbf{x} - \mathbf{p}) &= \sum_{i \in R_\mathbf{p}} \lambda_i (t - \rho) b_i \\ 0 &= \sum_{i \in R_\mathbf{p}} \lambda_i (t - \rho) b_i \\ 0 &= (t - \rho) \left(\sum_{i \in R_\mathbf{p}} \lambda_i b_i \right). \end{aligned}$$

But since the λ_i are non-negative (and not all of them are zero) and the b_i are positive, then it must be that $t = \rho$.

Thus, the system has only solutions of the form $(\mathbf{x}, \rho)$ and, hence, only the solution $(\mathbf{p}, \rho)$. This implies that the matrix $(A_{R_\mathbf{p}} | -b_{R_\mathbf{p}})$ has rank $d + 1$.

Choose R to be a basis for the rows of $(A_{R_\mathbf{p}} | b_{R_\mathbf{p}})$. □

A.2 Implementation considerations

Our implementation of Algorithm 2 uses the HiGHS MIP solver [9] to determine the feasibility of Integer Linear Problems defined in Proposition 11.

Since the MIP solver relies on numerical methods and hence is subject to numerical errors, we round all proposed solutions and check them for feasibility under exact linear algebra. We encountered no issues of this kind solving any of the ILPs needed to construct certificates for all volume vectors with volume at most 195.

In such cases, in lack of an exact MIP solver, a brute force approach could be used, checking all candidate translations within the bounding box of the zonotope.

Dyadic fundamental domains with small circumradius or small volume of convex hull can be obtained with the same algorithm, but turning the feasibility problem (3) from Proposition 11 into an optimization problem that minimizes some norm. This modification does not affect the search strategy or the types of voxels obtained in the final fundamental domain; it just gives the "best" representative of each type.

For example, the optimization problem for the Minkowski norm of P is particularly simple:

$$\begin{aligned} \text{minimize} \quad & \rho \in \mathbb{R} \\ \text{subject to} \quad & A\mathbf{x} - b\rho \leq -A\mathbf{c} - A_{\geq 0}\,\boldsymbol{\epsilon} \\ & \rho \geq 0 \\ & \mathbf{x} \in \mathbb{Z}^d. \end{aligned}$$

Optimizing with respect to this norm results in a dyadic fundamental domain fitting in the smallest possible dilation of the zonotope, among those with the types given by the breadth-first search.

The convex hull of the domain can be minimized even more by further subdividing all voxel types to reach a regular tree of any given depth, at the expense of solving many more ILPs.

Since our zonotopes are centrally symmetric around the origin, a voxel type will lie in our fundamental domain if and only if the opposite type does. Hence, we only need to check half of the voxels in the first subdivision of the unit cube, which automatically gives that we check only half of each level. This has the advantage of halving the execution time and producing centrally symmetric certificates, which are both smaller and visually clearer.

Implementations can easily avoid having to deal with rational matrices and vectors by scaling by the common denominator.