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

David Alcántara[†1], Francisco Criado[‡2], and Francisco Santos[§1]

[1]Departmento de Matemáticas, Estadística y Computación, Universidad de Cantabria, Santander, Spain
[2]Departamento de Matemáticas, CUNEF Universidad, Madrid, Spain

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, derived from the rephrasing of sLRC in terms of covering radii of certain zonotopes (Henze and Malikiosis 2017) plus an upper bound on the (integer) velocities that need to be checked (Malikiosis, Santos and Schymura, 2025).

As a tool for the proof we devise an algorithm for bounding the covering radius of lattice polytopes.

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 [13] in the language of diophantine approximation, and is currently proved up to $n = 6$ [1]. 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 [11] 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 [2].

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 $v_1, \dots, v_n$, and $s_1, \dots, s_n \in \mathbb{R}$ be real numbers, with the v_i distinct.*[1] *Then, there is a $t \in \mathbb{R}$ such that $\operatorname{dist}(v_i t + s_i, \mathbb{Z}) \geq \frac{1}{n+1}$ for every $i \in [n]$.*

This shifted version of the Lonely Runner Conjecture is only currently proved up to $n = 3$ ("four runners") [4, 12], and we prove it for five. 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 to be positive integers [5, 4]. We then rely on the following result of Malikiosis, Santos and Schymura:

Theorem 1 ([10, Corollary 1.15]) *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 1 *sLRC (Conjecture 1) holds for $n = 4$ (five runners).*

We also show there are only three primitive integer velocity vectors that are *tight*, in the sense that 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 $\operatorname{dist}(v_i t + s_i, \mathbb{Z}) \leq \frac{1}{n+1}$.

Theorem 3 *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 $s_1, \dots, s_4 \in \mathbb{R}$ such that, for all $t \in \mathbb{R}$, $\operatorname{dist}(v_i t + s_i, \mathbb{Z}) \leq \frac{1}{n+1}$, are $(1, 2, 3, 4)$, $(1, 3, 4, 6)$, and $(1, 3, 4, 7)$.*

*Research supported by grant PID2022-137283NB-C21, funded by MCIN/AEI/10.13039/501100011033.

†Funded by grant PRE2022-000020 of MCIN/AEI/10.13039/501100011033. Email: david.alcantara@unican.es.

‡Email: francisco.criado@cunef.edu.

§Email: francisco.santos@unican.es.

[1]In the original lonely runner conjecture dropping the condition that the v_i be distinct is no loss of generality, but the shifted version is false without it.

That the vector $(1,\dots,n)$ is tight for all $n\in\mathbb{N}$ is easy to show; and tightness of vector $(1,3,4,7)$ was already known to Wills for the non-shifted LRC conjecture [14]. Cusick and Pomerance [6] have shown that $(1,2,3,4)$ and $(1,3,4,7)$ are the only tight (primitive) instances of the non-shifted LRC for $n=4$. Obviously, tightness for the non-shifted version implies tightness for the shifted one. Our Theorem shows that the converse is not true.

Our method (as well as the proof of Theorem 1 in [10]) is based on the relation between the Lonely Runner conjecture (both shifted and original one) to $(n-1)$-zonotopes with n generators [7, 2]. 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 3 are computational; for each primitive velocity vector we build a fundamental domain of the integer lattice $\mathbb{Z}^3$ that fits in the (contracted) zonotope associated to it. This certifies that its covering radius satisfies the bound.

To prove tightness of the instances of Theorem 3, 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 [7, 2, 10].

A *zonotope* is any Minkowski sum of finitely many 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 one is $\mathbf{c}=-\frac{1}{2}\sum_{i=1}^{n}\mathbf{u}_i$. This 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.

2.1 Lonely runner zonotopes and volume vectors

Definition 4 *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*

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

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

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 [10]). Observe also that the generators and the volume vector satisfy

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

for some choice of signs. In fact, this equation (together with positivity) 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. That is, an element of $\mathrm{AGL}(n,\mathbb{Z}):=\mathbb{Z}^n\rtimes\mathrm{GL}(n,\mathbb{Z})$.

Proposition 1 ([10]) *For every integer vector $\mathbf{v}=(v_1,\dots,v_n)\in\mathbb{Z}^n_{>0}$ 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

As usual, 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. This includes all bounded full-dimensional polytopes.

Definition 5 (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 similar transforms of $\rho C+\mathbb{Z}^d$.

The zonotopal restatement of the *Shifted Lonely Runner Conjecture* is the following. Our statement is taken from [10] but the result is implicit in [7, 2, 4].

Proposition 2 ([7], see also [10, Proposition 1.8]) *Let $\mathbf{v}=(v_1,\dots,v_n)\in\mathbb{Z}^n_{>0}$ with pairwise distinct entries and $\gcd(v_1,\dots,v_n)=1$. 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. [4]), 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 6 *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 a fundamental domain.*

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., [9, 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. Our bound uses the concept of last covered point.

Definition 7 (Last covered point [3, 4]) *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 $\mathbf{p} \notin (\mu(C)C)^\circ + \mathbb{Z}^d$, where C° denotes the interior of C.*

The set of last covered points of a convex body C is always non-empty. Indeed, for every $\epsilon > 0$ there is a point $\mathbf{p}_\epsilon \notin (\rho - \epsilon)C + \mathbb{Z}^d$ and there is no loss of generality in assuming $\mathbf{p}_\epsilon \in \rho C$ for every ϵ, since $\rho C + \mathbb{Z}^d$ covers $\mathbb{R}^d$ and being in $(\rho - \epsilon)C + \mathbb{Z}^d$ is preserved under lattice translation. By compactness of ρC, there is a sequence $(\epsilon_i)_{i\in\mathbb{N}}$ converging to zero and with $\mathbf{p}_{\epsilon_i}$ converging to a point $\mathbf{p}$ of ρC. Now, $\mathbf{p}$ cannot be in $\rho' C + \mathbb{Z}^d$ for any $\rho' < \rho$, so $\mathbf{p}$ is a last covered point.

In the rest of the section, $P \subseteq \mathbb{R}^d$ 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 $i \in [m]$ or a 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 8 ([4, Lemma 3.1]) *Let $P = \{\mathbf{x} \in \mathbb{R}^d : Ax \le b\}$ 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*

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

has a unique solution which is a last covered point.

Proposition 3 *Let P be a rational polytope described by $Ax \le b$ 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 $\sqrt{\det\left((A|b)^T(A|b)\right)}$.*

Proof. The bound follows by applying Cramer's rule to the system (2) of Lemma 8, with $(\mathbf{x}, \rho)$ considered as variables, together with the following implication of Cauchy-Binet:

$$\det\left((A|b)^T(A|b)\right) = \sum_{R\in\binom{[m]}{d+1}} \det(A_R|b_R)^2 \ge \max_{R\in\binom{[m]}{d+1}} \det(A_R|b_R)^2. \qquad \square$$

Bounding the denominator of $\mu(P)$ is necessary in our algorithms, since it allows us to certify an exact upper bound from an approximate one, as follows.

Corollary 2 *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 Proposition 3). 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 $\mu(P) \le \rho - \frac{1}{sD}$. $\square$

3 Our Algorithms

In this section we describe the algorithm we have used to find fundamental domains within each sLR zonotope up to volume 195. See the appendix for a formal development.

Obtaining a representative zonotope for a velocity vector is discussed in [10]. We further simplify our representatives, reducing the length of their generators using the LLL algorithm. See details in Appendix.

3.1 Certifying an upper bound for the cov. radius

We here describe an algorithm to decide whether a facet-defined polytope $P = \{\mathbf{x} \in \mathbb{R}^d : A\mathbf{x} \le b\}$ contains a fundamental domain. By Lemma 6, 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 9 *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 one type are equivalent by integer translation, and voxel types are naturally arranged as an infinite rooted 2^ℓ-ary tree with the types of level ℓ at depth ℓ. We call this the infinite dyadic tree.

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.

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 all leaves fit inside, or the center of one leaf is found to have no translation inside, certifying a lower bound.

This algorithm is illustrated in Figure 1.

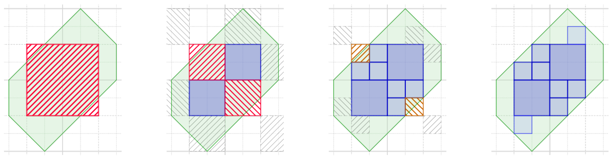

Figure 1: States of our Algorithm at different depths, applied to $\frac{1}{2}Z$ where Z is the 2-dimensional sLR zonotope with volume vector $(1, 2, 4)$.

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 4 *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:*

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

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

The search algorithm as described so far has two issues: on the one hand, if P does not contain a dyadic fundamental domain then the algorithm does not terminate; on the other hand, this may happen even if $\mu(P) \leq 1$, that is, if P contains a fundamental domain but not a dyadic one. If P is rational we can solve both issues thanks to Corollary 2.

Theorem 10 *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* $\mathbf{c} \in \frac{1}{2^\ell}\{0, \ldots, 2^\ell - 1\}^d$ *such that* $\left(\rho + \frac{1}{2sD}\right) P$ *does not intersect* $\mathbf{c} + \mathbb{Z}^d$*.*

See the appendix for a proof. 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.

References

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

[2] Matthias Beck, Serkan Hoşten, and Matthias Schymura. Lonely Runner Polyhedra. *Integers*, 19:#A29, 13 pp, 2019.

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

[4] Jana Cslovjecsek, Romanos Diogenes Malikiosis, Márton Naszódi, and Matthias Schymura. Computing the covering radius of a polytope with an application to lonely runners. *Combinatorica*, 42(4):463–490, February 2022.

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

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

[7] Matthias Henze and Romanos Diogenes 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.

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

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

[10] Romanos Diogenes Malikiosis, Francisco Santos, and Matthias Schymura. Linearly-exponential checking is enough for the lonely runner conjecture and some of its variants, 2024.

[11] Guillem Perarnau and Oriol Serra. The Lonely Runner Conjecture turns 60, 2024. `https://arxiv.org/abs/2409.20160`.

[12] Ludovic Rifford. On the time for a runner to get lonely, 2021.

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

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

A Algorithm details

A.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 $v = (v_1, \ldots, v_4) \in \mathbb{Z}^4$ 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 1, 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 5 *There are exactly* $2\,133\,561$ *vectors* $(v_1, v_2, v_3, v_4) \in \mathbb{Z}$ *with* $1 \le v_1 < v_2 < v_3 < v_4$, $\gcd(v_1, v_2, v_3, v_4) = 1$ *and* $\sum v_i \le 195$.

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 1.

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

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

Output : A matrix $M = (\mathbf{u}_1, \ldots, \mathbf{u}_n) \in \mathbb{Z}^{(n-1)\times n}$ such that $\mathbf{u}_1, \ldots, \mathbf{u}_n$ 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 $(v_1, \ldots, v_n)$. 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 columns of M'. 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 $(v_1, \ldots, v_n)$.

The generators obtained in this way typically have some large entries, resulting in 'long and skinny' zonotopes that are poorly conditioned for our method to compute covering radii. To overcome this we preprocess the generators in step 3, by performing an LLL lattice basis reduction to the rows of $B^{-1}M'$.[2] 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 $U = \{\mathbf{u}_1, \ldots, \mathbf{u}_n\}$ is done as follows, where we are identifying $\bigwedge^{d-1}\mathbb{R}^d \cong (\mathbb{R}^d)^*$ in the natural way.

Proposition 6 *Let* $Z = \frac{1}{2}\sum_{i=1}^n [-\mathbf{u}_i, \mathbf{u}_i]$ *be the* **0**-*symmetric zonotope with generators* $\mathbf{u}_1, \ldots, \mathbf{u}_n$. *Then*

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

where

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

A.2 Building a dyadic fundamental domain

In this section we present Algorithm 2, the concrete algorithm that explores the infinite dyadic tree to decide 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 6.

The termination of this algorithm follows from Theorem 10, of which we give now proof.

Proof. [of Theorem 10] For part (1) we only need to use that, for $P^+ = \left(\rho + \frac{1}{2sD}\right)$ as in the algorithm,

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

For each $\ell \in \mathbb{N}$ let D_ℓ be the union of all the dyadic voxels 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.

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

For part (2) we use that $\mu(P) > \rho$ implies (by Corollary 2) 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 6, which asserts the existence of an open set $W \subset \mathbb{R}^d \setminus (P^+ + \mathbb{Z}^d)$. ☐

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 10.

1 Let D be a bound on the denominator of $\mu(P)$, such as the one from Proposition 3.

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 minimum depth.

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

Theorem 11 *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 2 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 6 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 10 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 $\leq \ell$ has a representative contained in P^+, so the algorithm will never enter the "else" in line 14 with a voxel of depth $\leq \ell$. Hence, the algorithm can perform the while loop only finitely many times before N becomes empty.

If $\mu(P) > \rho$, Theorem 10 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). ☐

A.3 Implementation considerations

Our implementation of Algorithm 2 uses the HiGHS MIP solver [8] to determine the feasibility of Integer Linear Problems defined in Proposition 4.

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.

If *pretty* dyadic fundamental domains are desired, the feasibility problem from Proposition 4 can be turned into an optimization problem, minimizing some norm, such as the Minkowski norm of P, which results in a dyadic fundamental domain that fits in the smallest contraction of the zonotope among those whose leaves have equal depth.

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

Since we assume our zonotopes to be centrally symmetric around the origin, our implementation also considers the centered unit cube, and avoids checking half of the voxels in its first subdivision, building the other half by central symmetry.