

The irreducible vectors of a lattice: Some theory and applications

Emmanouil Doulgerakis, Thijs Laarhoven, and Benne de Weger

Eindhoven University of Technology Eindhoven, The Netherlands
{e.doulgerakis,b.m.m.d.weger}@tue.nl, mail@thijs.com

Abstract. The main idea behind lattice sieving algorithms is to reduce a sufficiently large number of lattice vectors with each other so that a set of short enough vectors is obtained. It is therefore natural to study vectors which cannot be reduced. In this work we give a concrete definition of an irreducible vector and study the properties of the set of all such vectors. We show that the set of irreducible vectors is a subset of the set of relevant vectors and study its properties. For extremal lattices this set may contain as many as 2^n vectors, which leads us to define the notion of a complete system of irreducible vectors, whose size can be upper-bounded by the kissing number. We study properties of this set and observe a close relation to heuristic sieving algorithms. Finally we briefly examine the use of this set in the study of lattice problems such as SVP, SIVP and CVPP. The introduced notions, as well as various results derived along the way, may provide further insights into lattice algorithms and motivate new research into understanding these algorithms better.

Keywords: lattices · relevant vectors · irreducible vectors · sieving algorithms

1 Introduction

The need for quantum-resistant cryptography has led to rapid developments in the area of lattice-based cryptography, mainly spurred by the NIST PQ-Crypto competition. Large scale deployment of lattice-based cryptosystems in the near future becomes realistic. This continues to make the deeper understanding of lattice problems an urgent research topic.

In 2010 Micciancio and Voulgaris, based also on previous work [1], described deterministic $\tilde{O}(2^{2n})$ -time and $\tilde{O}(2^n)$ -space algorithms to solve some of the most important lattice problems (such as SVP, SIVP and CVP) [22] in dimension n . This result mainly relies on an algorithm to compute the set of relevant vectors of (the Voronoi cell of) a lattice. Even though this is a very interesting result, the constants in the exponents of time and space complexities of the Micciancio–Voulgaris algorithm make it impractical, even for moderate dimensions.

The set of relevant vectors was first introduced in 1908 by Voronoi [35]. It provides a useful representation of the Voronoi cell of a lattice. Even though the

set of relevant vectors seems to hold the key for solving many lattice problems, its expected size makes it impractical. This becomes even more clear when that size is compared to the (time and) space complexity of algorithms used in practice for solving lattice problems such as [2, 6, 13].

In this work, we introduce a different set of lattice vectors, which appears to serve as a bridge between the provable results relying on the set of relevant vectors and heuristic sieving algorithms [3, 23, 27].

Notions of irreducibility are considered to be fundamental in many areas. Often irreducibility is defined with respect to multiplication. Since a lattice is an additive object, we will however use an additive notion of irreducibility. Clearly the notion of *lattice basis* could be seen as such a construct, but it has been observed to be a too weak notion to provide, on its own, interesting results for lattice problems. Our new notion of *irreducible vectors* provides us with a set of lattice vectors, larger than a basis but smaller than the set of relevant vectors, and possessing interesting properties. To the best of our knowledge this definition is new in the area of lattices.

Contributions. In this paper we define a notion of irreducibility for a lattice vector. As a first result we show that every irreducible vector of a lattice belongs to the lattice’s set of relevant vectors. Hence, the set of irreducible vectors which we denote by $\text{Irr}(\mathcal{L})$ is finite. Additionally, it is shown that the set of irreducible vectors generates the lattice and also contains vectors achieving all the successive minima of the lattice. Finally, the set of irreducible vectors of the root lattices A_n , D_n and their duals A_n^* , D_n^* is examined as they prove to be interesting extreme cases.

As it turns out, the set $\text{Irr}(\mathcal{L})$ can be as big as the set of relevant vectors. In order to get a set of cardinality provably smaller than 2^n , a complete system of irreducible vectors is defined, which is denoted by $\text{P}(\mathcal{L})$. This set inherits the aforementioned properties of the set $\text{Irr}(\mathcal{L})$ and also it is proved that $|\text{P}(\mathcal{L})| < 2^{0.402n}$ where n is the rank of the lattice. Heuristically it is expected that $\text{P}(\mathcal{L})$ will have a cardinality of $2^{0.21n}$. From a computational point of view, it is shown that slightly modified versions of already existing sieving algorithms asymptotically converge to such a set (modulo sign). This statement is further supported by experimental results. Finally, we discuss the applicability of $\text{P}(\mathcal{L})$ in showing that sieving algorithms like the ones described in [3, 23] can be used for tackling SVP, SIVP and computing the kissing number of a lattice. Additionally we discuss the applicability of $\text{P}(\mathcal{L})$ as preprocessing data in a CVPP algorithm which we call “the tuple slicer”. The tuple slicer can provide a time–memory trade-off without the use of rerandomisations.

Outline. The rest of the paper is organised as follows. In Section 2 we introduce notation and give some background about lattices. Section 3 includes some prior work on the set of relevant vectors. The definition of irreducible vectors is given in Section 4 along with the first results regarding this new notion. In Section 5 we mention theoretical as well as experimental results on computing a complete system of irreducible vectors. Section 6 provides some initial arguments about

the link between the new notions defined and the study of lattice algorithms and problems. Finally, Section 7 concludes with a discussion of potentially interesting future research topics.

2 Preliminaries

To fix notation, let $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\} \subset \mathbb{R}^n$ be a set of linearly independent vectors, which we may also interpret as a matrix with columns \mathbf{b}_i . The lattice generated by \mathbf{B} is defined as $\mathcal{L} = \mathcal{L}(\mathbf{B}) := \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^n\}$. In this paper we deal with full rank lattices unless indicated otherwise. We assume that the reader is familiar with notions such as the *volume* $\text{Vol}(\mathcal{L}) := |\det(\mathbf{B})|$, the *successive minima* $\lambda_i(\mathcal{L}) := \min\{\max_i \|\mathbf{x}_i\| \mid \mathbf{x}_1, \dots, \mathbf{x}_i \in \mathcal{L} \text{ are linearly independent}\}$, in particular the *first successive minimum* $\lambda_1(\mathcal{L}) = \min_{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\}} \|\mathbf{v}\|$. We refer to [21] for further details on these basic notions.

Definition 1 (First shell). *Let \mathcal{L} be a lattice. We define*

$$S_1(\mathcal{L}) := \{\mathbf{v} \in \mathcal{L} \mid \|\mathbf{v}\| = \lambda_1(\mathcal{L})\}. \quad (1)$$

We call $S_1(\mathcal{L})$ the first shell of \mathcal{L} .

The following two well known concepts will be of major importance for our work, so we define them explicitly.

Definition 2 (Voronoi cell). *The Voronoi cell $\mathcal{V}(\mathcal{L})$ of a full rank lattice \mathcal{L} is the set of points in \mathbb{R}^n which are closer to the origin than to any other lattice point, i.e.*

$$\mathcal{V}(\mathcal{L}) := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \leq \|\mathbf{x} - \mathbf{v}\| \forall \mathbf{v} \in \mathcal{L}\}. \quad (2)$$

If it's clear what \mathcal{L} is, we may use \mathcal{V} instead of $\mathcal{V}(\mathcal{L})$. Closely related to the Voronoi cell of the lattice is the set of relevant vectors.

Definition 3 (Relevant vectors). *The set of relevant vectors $R(\mathcal{L})$ is*

$$R(\mathcal{L}) := \{\mathbf{r} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid (\mathbf{r} + \mathcal{V}) \text{ and } \mathcal{V} \text{ share a non-empty boundary}\}. \quad (3)$$

Let $\mathcal{B}(\mathbf{x}, r) := \{\mathbf{y} \in \mathbb{R}^n \mid \|\mathbf{y} - \mathbf{x}\| \leq r\}$ denote the closed n -dimensional ball with center \mathbf{x} and radius r . Finally we have the *kissing number* τ_n , defined as the maximum number of equal n -dimensional spheres that can be made to touch another central sphere of the same size without intersecting.

See [21] for an overview of the main hard lattice problems that we will consider in this paper, namely the Shortest Vector Problem (SVP), determining the kissing number, the Shortest Independent Vector Problem (SIVP), and the Closest Vector Problem (CVP) and its Preprocessing variant (CVPP).

3 Previous Work

In this section we give an overview of known results on the set of relevant vectors. This is done for a matter of completeness but also in order to indicate what kind of results we would like to obtain for the set of irreducible vectors which we will define later.

For $\mathbf{v} \in \mathcal{L}$ we define $H(\mathbf{v}) := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \leq \|\mathbf{x} - \mathbf{v}\|\}$, to relate the Voronoi cell of a lattice to its relevant vectors.

Proposition 1 (Relevant vectors). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . The set of relevant vectors $R(\mathcal{L})$ is the minimal set $L \subset \mathcal{L}$ such that*

$$\mathcal{V}(\mathcal{L}) = \bigcap_{\mathbf{v} \in L} H(\mathbf{v}). \quad (4)$$

In order to get a more practical description of the relevant vectors the following theorem is used.

Theorem 1 (Identifying relevant vectors [35]). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\}$. Then $\mathbf{v} \in R(\mathcal{L})$ if and only if $\mathbf{0}$ and \mathbf{v} are the only closest vectors of \mathcal{L} to $\frac{1}{2}\mathbf{v}$.*

This implies that

$$R(\mathcal{L}) = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \|\frac{1}{2}\mathbf{v} - \mathbf{x}\| > \|\frac{1}{2}\mathbf{v}\| \quad \forall \mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}, \mathbf{v}\}\} \quad (5)$$

$$= \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \langle \mathbf{v}, \mathbf{x} \rangle < \|\mathbf{x}\|^2 \quad \forall \mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}, \mathbf{v}\}\} \quad (6)$$

Remark 1. It holds that $\mathbf{0} \notin R(\mathcal{L})$. Also note that if $\mathbf{v} \in R(\mathcal{L})$ then $-\mathbf{v} \in R(\mathcal{L})$.

Remark 2. The condition $\langle \mathbf{v}, \mathbf{x} \rangle < \|\mathbf{x}\|^2$ needs to be checked only for $\mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}\}$ such that $\|\mathbf{x}\| < \|\mathbf{v}\|$, because otherwise it is trivially true.

For checking if a vector is relevant, the following lemma is useful.

Lemma 1 (Identifying non-relevant vectors [22]). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n , and $\mathbf{v} \in \mathcal{L}$. If $\mathbf{v} \notin R(\mathcal{L})$ then there exists $\mathbf{r} \in R(\mathcal{L})$ such that $\langle \mathbf{v}, \mathbf{r} \rangle \geq \|\mathbf{r}\|^2$.*

Also a lower bound for the set $R(\mathcal{L})$ can be obtained by the following trivial lemma.

Lemma 2 (All shortest vectors are relevant). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . It holds that $S_1(\mathcal{L}) \subseteq R(\mathcal{L})$.*

Equality in the above lemma holds for a very special type of lattices.

Theorem 2 (Root lattices [28]). *$S_1(\mathcal{L}) = R(\mathcal{L})$ iff \mathcal{L} is a root lattice.*

The following theorem by Minkowski gives an upper bound on the size of $R(\mathcal{L})$.

Theorem 3 (Upper bound on $|\mathbf{R}(\mathcal{L})|$ [24]). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . It holds that $|\mathbf{R}(\mathcal{L})| \leq 2(2^n - 1)$.*

Apart from an upper bound we can also obtain a lower bound on $|\mathbf{R}(\mathcal{L})|$.

Remark 3. For the lattice \mathbb{Z}^n is true that $|\mathbf{R}(\mathbb{Z}^n)| = 2n$ (see [8]). As the set $\mathbf{R}(\mathcal{L})$ needs to have n linearly independent vectors in order the volume of $\mathcal{V}(\mathcal{L})$ to be finite then $2n \leq |\mathbf{R}(\mathcal{L})|$ is a tight lower bound.

Proposition 2 (Volume of the Voronoi cell). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . It holds that $\text{Vol}(\mathcal{L}) = \text{Vol}(\mathcal{V}(\mathcal{L}))$.*

Proposition 3 (Relevant vectors generate). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . There exists a generating set \mathbf{G} of \mathcal{L} such that $\mathbf{G} \subseteq \mathbf{R}(\mathcal{L})$.*

Proof. Every vector $\mathbf{v} \in \mathcal{L}$ can be “reduced” to the $\mathbf{0}$ -vector through reductions by elements of $\mathbf{R}(\mathcal{L})$. Thus, $\mathbf{R}(\mathcal{L})$ spans the entire lattice. \square

So far we have mentioned a number of properties and definitions on the relevant vectors of a lattice. Computing them is however a different matter. The following result is the current state of the art on this.

Theorem 4 (Finding all relevant vectors [22]). *There exists a deterministic $\tilde{O}(2^{2n})$ -time and $\tilde{O}(2^n)$ -space algorithm which, given an n -rank lattice \mathcal{L} with basis \mathbf{B} , outputs the set of relevant vectors.*

4 Irreducibility of lattice vectors

4.1 The set of irreducible vectors

Inspired by number theoretic notions of (multiplicative) irreducibility, we introduce a similar concept for lattice (additively structured).

Definition 4 (Irreducibility). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\}$. The vector \mathbf{v} is called k -irreducible iff $\nexists \mathbf{v}_1, \dots, \mathbf{v}_k \in \mathcal{L}$ such that $\|\mathbf{v}_i\| < \|\mathbf{v}\|$ and $\mathbf{v}_1 + \dots + \mathbf{v}_k = \mathbf{v}$. For the special case $k = 2$, \mathbf{v} will be just called irreducible.*

Remark 4. The definition of k -irreducible vectors implies that if a vector is k -irreducible then it is also $(k - 1)$ -irreducible. This observation allows the construction of a chain of subsets based on the notion of irreducibility.

In this work we are going to focus on the properties of 2-irreducibility. Further research on the notion of k -irreducibility for $k > 2$ is left for future research.

Definition 5 (Set of irreducible vectors). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . We define*

$$\text{Irr}(\mathcal{L}) := \{\mathbf{v} \in \mathcal{L} \mid \mathbf{v} \text{ is irreducible}\}. \quad (7)$$

Remark 5. It holds that $\mathbf{0} \notin \text{Irr}(\mathcal{L})$. Also, if $\mathbf{v} \in \text{Irr}(\mathcal{L})$ then $-\mathbf{v} \in \text{Irr}(\mathcal{L})$.

The above properties hold for the set of relevant vectors as well and this is not a coincidence as we will see. First we show that this set is not empty, and indeed that it also contains vectors achieving the first successive minimum.

Lemma 3 (Shortest vectors are irreducible). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . It holds that:*

$$S_1(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L}). \quad (8)$$

Proof. Let $\mathbf{v} \in S_1(\mathcal{L})$. Then clearly $\mathbf{v} \neq \mathbf{0}$. Assume that $\mathbf{v} \notin \text{Irr}(\mathcal{L})$, so there exist $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{L}$ such that $\|\mathbf{v}_i\| < \|\mathbf{v}\|$ and $\mathbf{v}_1 + \mathbf{v}_2 = \mathbf{v}$. As $\mathbf{v} \in S_1(\mathcal{L})$ this implies that $\|\mathbf{v}_i\| < \lambda_1(\mathcal{L})$ and thus $\|\mathbf{v}_i\| = 0$. Hence, we get $\mathbf{v}_1 = \mathbf{v}_2 = \mathbf{0}$, which contradicts $\mathbf{v} \neq \mathbf{0}$. \square

Remark 6. It can be easily checked that Lemma 3 would still hold under the notion of k -irreducibility for $k > 2$. Therefore we can conclude that k -irreducibility is not leading to a trivially empty set of vectors for $k > 2$. One may expect that it will also include a lattice basis.

We show that something similar occurs for the rest of the successive minima as well.

Definition 6 (Sublattice spanned by short vectors). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $1 \leq i \leq n$. We define \mathcal{L}_λ to be the sublattice spanned by all the vectors in \mathcal{L} with norm strictly less than λ .*

Proposition 4 (Identifying irreducible vectors). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n , and $\mathbf{v} \in \mathcal{L}$ satisfying $\|\mathbf{v}\| = \lambda_i := \lambda_i(\mathcal{L})$ for some $1 \leq i \leq n$. If $\mathbf{v} \notin \mathcal{L}_{\lambda_i}$ then \mathbf{v} is irreducible.*

Proof. It has been already proven in Lemma 3 that this is true for $i = 1$ so we can consider $i \geq 2$. Assume $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\| = \lambda_i(\mathcal{L})$ for some $2 \leq i \leq n$, $\mathbf{v} \notin \mathcal{L}_{\lambda_i}$ and \mathbf{v} is not irreducible. Then there exist $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{L}$ such that $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$ and $\|\mathbf{v}_j\| < \|\mathbf{v}\|$ for $j = 1, 2$. Clearly $\mathbf{v}_j \neq \mathbf{0}$. As $\|\mathbf{v}_j\| < \|\mathbf{v}\| = \lambda_i(\mathcal{L})$ this implies that $\mathbf{v}_j \in \mathcal{L}_{\lambda_i}$ for $j = 1, 2$. This further implies that $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2 \in \mathcal{L}_{\lambda_i}$, contradiction. \square

Remark 7. Proposition 4 points out that a lattice vector achieving a successive minimum is not necessarily irreducible. An enlightening example of such an occasion is the following. Consider the lattice $\mathcal{L} = \mathcal{L}(\mathbf{B})$ generated by the matrix

$$\mathbf{B} = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 10 \end{pmatrix}. \quad (9)$$

Then $\lambda_1(\mathcal{L}) = 3$, $\lambda_2(\mathcal{L}) = 4$ and $\lambda_3(\mathcal{L}) = 10$. The vector $\mathbf{v} = (6, 8, 0)$ is such that $\|\mathbf{v}\| = \lambda_3(\mathcal{L})$ but \mathbf{v} is not irreducible as it can be written as a sum of shorter vectors i.e. $\mathbf{v} = (6, 0, 0) + (0, 8, 0)$. The reason why \mathbf{v} fails to be irreducible is that it belongs to the sublattice \mathcal{L}_{λ_3} .

Corollary 1 (Irreducible vectors and successive minima). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . For every $i = 1, \dots, n$ there exists a vector $\mathbf{v} \in \text{Irr}(\mathcal{L})$ such that $\|\mathbf{v}\| = \lambda_i(\mathcal{L})$.*

Proof. By Proposition 4 it suffices to show that for every $i = 1, \dots, n$ there exists a vector $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\| = \lambda_i(\mathcal{L})$ and $\mathbf{v} \notin \mathcal{L}_{\lambda_i}$. Assume that for every vector $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\| = \lambda_i(\mathcal{L})$ for some fixed $2 \leq i \leq n$ it holds $\mathbf{v} \in \mathcal{L}_{\lambda_i}$. For convenience we define $\lambda_0(\mathcal{L}) = 0$. Let k be $\min_{1 \leq j \leq i} j$ such that $\lambda_j(\mathcal{L}) = \lambda_i(\mathcal{L})$ and therefore $\lambda_{k-1}(\mathcal{L}) < \lambda_k(\mathcal{L}) = \lambda_i(\mathcal{L})$. Then \mathcal{L}_{λ_i} has rank $k-1$ as $\lambda_{k-1}(\mathcal{L}) < \lambda_k(\mathcal{L})$. If $k-1 = 0$ then we are done as this would imply $\mathbf{v} = \mathbf{0}$. If $k-1 \geq 1$ then \mathbf{v} belongs to the sublattice containing all the shorter vectors than it, \mathcal{L}_{λ_i} and this sublattice is of rank $k-1$. Thus any choice of $k-1$ vectors such that $\max\{\|\mathbf{v}_1\|, \dots, \|\mathbf{v}_{k-1}\|, \|\mathbf{v}\|\} = \|\mathbf{v}\|$ will result in a linearly dependent set. Hence it cannot be that $\lambda_k(\mathcal{L}) = \|\mathbf{v}\|$, contradiction. \square

Apart from vectors reaching the successive minima, it can be shown that the set $\text{Irr}(\mathcal{L})$ contains a generating set of the lattice as well.

Proposition 5 (Irreducible vectors generate). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . There exists a generating set \mathbf{G} of \mathcal{L} such that $\mathbf{G} \subseteq \text{Irr}(\mathcal{L})$.*

Proof. We will prove that the set $\text{Irr}(\mathcal{L})$ spans the lattice and therefore it includes a generating set. Let $\mathbf{v} \in \mathcal{L}$. If $\nexists \mathbf{v}_1, \mathbf{v}_2 \in \mathcal{L}$ with $\|\mathbf{v}_i\| < \|\mathbf{v}\|$ such that $\mathbf{v}_1 + \mathbf{v}_2 = \mathbf{v}$ then $\mathbf{v} \in \text{Irr}(\mathcal{L})$. If there exist such \mathbf{v}_i then write $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$. If the $\mathbf{v}_i \in \text{Irr}(\mathcal{L})$ then we are done. If not then further reduce the vectors \mathbf{v}_i such that they are written as a sum of two strictly shorter vectors. As in each step the length of the vectors strictly reduces and there is a finite number of lattice points in $\mathcal{B}(\mathbf{0}, \|\mathbf{v}\|)$, after a finite number of steps we will reach a state where $\mathbf{v} = \sum \mathbf{p}_i$ and $\mathbf{p}_i \in \text{Irr}(\mathcal{L})$. This concludes the proof. \square

Given the result of Proposition 5 the following conjecture can be formulated.

Conjecture 1. Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . The set $\text{Irr}(\mathcal{L})$ contains a basis of \mathcal{L} .

Our next goal is to derive some more explicit descriptions of the set $\text{Irr}(\mathcal{L})$.

Lemma 4 (Characterizing the set of irreducible vectors). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . It holds that*

$$\text{Irr}(\mathcal{L}) = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \forall \mathbf{x} \in \mathcal{L} \text{ with } \|\mathbf{x}\| < \|\mathbf{v}\| \text{ it holds } \|\mathbf{v} - \mathbf{x}\| \geq \|\mathbf{v}\|\} \quad (10)$$

$$= \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \forall \mathbf{x} \in \mathcal{L} \text{ with } \|\mathbf{x}\| < \|\mathbf{v}\| \text{ it holds } 2\langle \mathbf{v}, \mathbf{x} \rangle \leq \|\mathbf{x}\|^2\}. \quad (11)$$

Proof. Let $A = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \forall \mathbf{x} \in \mathcal{L} \text{ with } \|\mathbf{x}\| < \|\mathbf{v}\| \text{ it holds } \|\mathbf{v} - \mathbf{x}\| \geq \|\mathbf{v}\|\}$. Let $\mathbf{p} \in \text{Irr}(\mathcal{L})$ and $\mathbf{v} \in \mathcal{L}$ with $\|\mathbf{v}\| < \|\mathbf{p}\|$. Then as $\mathbf{p} \in \text{Irr}(\mathcal{L})$ we get $\|\mathbf{p} - \mathbf{v}\| \geq \|\mathbf{p}\|$ because otherwise \mathbf{p} would have a decomposition into two shorter vectors, thus $\mathbf{p} \in A$. This gives $\text{Irr}(\mathcal{L}) \subseteq A$. Next, let $\mathbf{v} \in A$, and write $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$ for some $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{L}$. If $\|\mathbf{v}_1\| < \|\mathbf{v}\|$ then as $\mathbf{v} \in A$ we get $\|\mathbf{v} - \mathbf{v}_1\| \geq \|\mathbf{v}\|$ and hence we do not get a decomposition of \mathbf{v} in two shorter vectors. If $\|\mathbf{v}_1\| \geq \|\mathbf{v}\|$ this is trivially true. Thus $\mathbf{v} \in \text{Irr}(\mathcal{L})$. This implies equality (10) and equality (11) is an immediate consequence. This concludes the proof. \square

Even though this lemma is rather straightforward it implies an interesting result for the set $\text{Irr}(\mathcal{L})$.

Proposition 6 (Irreducible vectors are relevant). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . Every irreducible vector of \mathcal{L} is also a relevant vector of \mathcal{L} , hence:*

$$\text{Irr}(\mathcal{L}) \subseteq \text{R}(\mathcal{L}). \quad (12)$$

Proof. As we already saw by Theorem 1, we can write the set $\text{R}(\mathcal{L})$ as $\text{R}(\mathcal{L}) = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \langle \mathbf{v}, \mathbf{x} \rangle < \|\mathbf{x}\|^2 \ \forall \mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}, \mathbf{v}\}\}$ and we can further improve that description to

$$\text{R}(\mathcal{L}) = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \forall \mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}\} \text{ with } \|\mathbf{x}\| < \|\mathbf{v}\| \text{ it holds } \langle \mathbf{v}, \mathbf{x} \rangle < \|\mathbf{x}\|^2\}.$$

For the set of irreducible vectors we got from Lemma 4 that

$$\text{Irr}(\mathcal{L}) = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \forall \mathbf{x} \in \mathcal{L} \text{ with } \|\mathbf{x}\| < \|\mathbf{v}\| \text{ it holds } 2\langle \mathbf{v}, \mathbf{x} \rangle \leq \|\mathbf{x}\|^2\}.$$

Thus by carefully checking these two descriptions for the sets $\text{R}(\mathcal{L})$ and $\text{Irr}(\mathcal{L})$ it suffices to prove that if $\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\}$ and $\mathbf{x} \in \mathcal{L} \setminus \{\mathbf{0}\}$ with $\|\mathbf{x}\| < \|\mathbf{v}\|$ then $2\langle \mathbf{v}, \mathbf{x} \rangle \leq \|\mathbf{x}\|^2 \Rightarrow \langle \mathbf{v}, \mathbf{x} \rangle < \|\mathbf{x}\|^2$.

If $\langle \mathbf{v}, \mathbf{x} \rangle \leq 0$ this is trivially true as $\mathbf{x} \neq \mathbf{0}$. Also if $\langle \mathbf{v}, \mathbf{x} \rangle > 0$ then $\langle \mathbf{v}, \mathbf{x} \rangle < 2\langle \mathbf{v}, \mathbf{x} \rangle$ and the result follows. \square

Remark 8. Combining the result of Lemma 3 and Proposition 6 we get that $S_1(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L}) \subseteq \text{R}(\mathcal{L})$. Therefore $\text{Irr}(\mathcal{L})$ is finite.

We already saw that for the case of root lattices it holds $S_1(\mathcal{L}) = \text{R}(\mathcal{L})$. This implies that for the root lattices it also holds that $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L}) = \text{R}(\mathcal{L})$. Thus, the sets $S_1(\mathcal{L})$ and $\text{R}(\mathcal{L})$ are tight inclusions of $\text{Irr}(\mathcal{L})$.

We expect that in general though it will hold $S_1(\mathcal{L}) \subsetneq \text{Irr}(\mathcal{L}) \subsetneq \text{R}(\mathcal{L})$. A question that might be of interest is when and if $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L}) \subsetneq \text{R}(\mathcal{L})$ or $S_1(\mathcal{L}) \subsetneq \text{Irr}(\mathcal{L}) = \text{R}(\mathcal{L})$ are possible.

We believe that lattices satisfying either of these properties will be very special and highly symmetric. The reason why we believe this, is that some already well known very special families of lattices satisfy these properties. Namely, in Appendix A we will prove the following two theorems.

Theorem 5 (The root lattices D_n^*). *Let $n \in \mathbb{N}$ with $n \geq 5$. Then for the lattice D_n^* it holds that $S_1(D_n^*) \subsetneq \text{Irr}(D_n^*) = \text{R}(D_n^*)$. Furthermore $|\text{Irr}(D_n^*)| = 2^n + 2n$.*

Theorem 6 (The root lattices A_n^*). *Let $n \in \mathbb{N}$ with $n \geq 3$. Then for the lattice A_n^* it holds that $S_1(A_n^*) = \text{Irr}(A_n^*) \subsetneq \text{R}(A_n^*)$. Furthermore $|\text{Irr}(A_n^*)| = 2(n+1)$.*

Additionally the famous Leech lattice A_{24} [8, p. 131] satisfies the property $S_1(A_{24}) = \text{Irr}(A_{24}) \subsetneq \text{R}(A_{24})$. We will actually be able to prove in the next subsection that for every lattice that reaches the kissing number τ_n it holds that $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L})$.

4.2 A complete system of irreducible vectors

The special family of lattices D_n^* indicates that the set $\text{Irr}(\mathcal{L})$ can become as big as $R(\mathcal{L})$ and actually grow as much in size as 2^n . However, our goal is to obtain a subset of $\text{Irr}(\mathcal{L})$ which is closely related to it but also provably smaller than 2^n .

Definition 7 (Equivalence relation on $\text{Irr}(\mathcal{L})$). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . We define an equivalence relation in $\text{Irr}(\mathcal{L})$ in the following way.*

$$\text{Let } \mathbf{v}_1, \mathbf{v}_2 \in \text{Irr}(\mathcal{L}) \text{ then } \mathbf{v}_1 \sim \mathbf{v}_2 \text{ iff } \|\mathbf{v}_1\| = \|\mathbf{v}_2\|. \quad (13)$$

From each equivalence class we will consider at least two representatives. We choose them in the following way and we will explain afterwards why.

Definition 8 (Representatives of equivalence classes). *For each equivalence class $S = \{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ of $\text{Irr}(\mathcal{L})$ according to (13) we choose a subset $\tilde{S} \subseteq S$ such that the following two conditions hold:*

- (i) *If $\mathbf{v} \in \tilde{S}$ then also $-\mathbf{v} \in \tilde{S}$.*
- (ii) *\tilde{S} is a maximal subset of S such that for every pair of vectors $\mathbf{v}_1, \mathbf{v}_2 \in \tilde{S}$ with $\mathbf{v}_2 \neq -\mathbf{v}_1$ it holds that $\|\mathbf{v}_1 + \mathbf{v}_2\| \geq \|\mathbf{v}_1\|$.*

The main motivation is that the new set of vectors which will be built under these rules will include irreducible vectors whose pairwise angle is “big” as we will prove later. However, there are several details of this definition which should be clarified. First of all, from the definition it follows that for an equivalence class we consider at least two representatives, which is not usually done. The reasons for this are the following.

Initially, for the subset of $\text{Irr}(\mathcal{L})$ which we are trying to define, we would like it to inherit the property of $\text{Irr}(\mathcal{L})$ that if \mathbf{v} belongs to it then also $-\mathbf{v}$ belongs to it. A second, more important reason is that choosing only one representative per equivalence class could lead to a set that does not even span the lattice (for example in the case of root lattices, or whenever $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L})$).

The second condition of the definition implies that for every element \mathbf{v} of a class S which is not included in \tilde{S} there exists a vector $\tilde{\mathbf{v}} \in \mathcal{L}$ such that $\|\mathbf{v} - \tilde{\mathbf{v}}\| < \|\mathbf{v}\|$. From this point of view the remaining elements of a class S which are not included in \tilde{S} can be generated by the elements of \tilde{S} plus some strictly shorter vector. In order to ensure that this holds we take \tilde{S} to be maximal. Also by taking \tilde{S} to be maximal we make sure that the set \tilde{S} contains as much information about the class as possible.

Remark 9. Choosing a representative set \tilde{S} of a class S can be translated into a graph problem. We define a graph where the set of vertices is the equivalence class, and there exists an edge between two vertices iff the difference of the corresponding vectors is strictly shorter than both of them. Then choosing a set of representatives translates to finding a maximal subset of vertices that are not adjacent, while keeping the symmetry about $\mathbf{0}$. In terms of graph theory this can be phrased as finding a special independent set of the graph. This idea is further analysed in Appendix B.

Definition 9 (Complete system of irreducible vectors). Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . We define a complete system of irreducible vectors of \mathcal{L} to be:

$$P(\mathcal{L}) := \bigcup_{S \in \text{Irr}(\mathcal{L})/\sim} \tilde{S}. \quad (14)$$

Remark 10. By the definition of $P(\mathcal{L})$ it follows that it is not uniquely determined. Its contents depend on the choice of representatives \tilde{S} for each $S \in \text{Irr}(\mathcal{L})/\sim$. Furthermore, this also implies that the size of $P(\mathcal{L})$ can vary.

Remark 11. By the fact that for each class S of $\text{Irr}(\mathcal{L})/\sim$ we have $\tilde{S} \subseteq S$ we get that $P(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L})$. Also the class of $\text{Irr}(\mathcal{L})/\sim$ containing all the shortest vectors i.e. $S_1(\mathcal{L})$ will be entirely included in $P(\mathcal{L})$ as any pairwise sum of vectors (for non-trivial pairs) in this class will be longer or equally long by definition. Thus we can conclude that

$$S_1(\mathcal{L}) \subseteq P(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L}) \subseteq R(\mathcal{L}).$$

We will also give an example in order to illustrate this definition.

Example 1. Let $\mathcal{L} = \mathcal{L}(\mathbf{B})$ be the lattice generated by the columns of the matrix

$$\mathbf{B} = \begin{pmatrix} 0 & 0 & 0 & -1 & 3 \\ -1 & -1 & 0 & 1 & 0 \\ 0 & 0 & -2 & 1 & 0 \\ 1 & -1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 2 & 1 \end{pmatrix}.$$

We find the sets $S_1(\mathcal{L}), P(\mathcal{L}), \text{Irr}(\mathcal{L}), R(\mathcal{L})$.

In fact \mathbf{B} is an LLL-reduced basis [19] of the lattice. By means of enumeration one could verify that $S_1(\mathcal{L}) = \{\pm(0, -1, 0, 1, 0)\}$. By running an algorithm that computes the set of relevant vectors like [34] in SAGE [33] we get

$$\begin{aligned} R(\mathcal{L}) = \{ & \pm(0, -1, 0, 1, 0), \\ & \pm(0, -1, 0, -1, 1), \\ & \pm(0, 0, 2, 0, 0), \\ & \pm(-1, 0, -1, 1, 2), \pm(-1, 0, 1, 1, 2), \pm(-1, 1, -1, 0, 2), \pm(-1, 1, 1, 0, 2), \\ & \pm(-1, 1, -1, 2, 1), \pm(-1, 1, 1, 2, 1), \pm(-1, 2, -1, 1, 1), \pm(-1, 2, 1, 1, 1), \\ & \pm(3, 1, 0, 0, 1), \pm(3, 0, 0, 1, 1), \\ & \pm(2, 1, -1, 1, 3), \pm(2, 1, 1, 1, 3), \\ & \pm(2, 2, -1, 2, 2), \pm(2, 2, 1, 2, 2)\}. \end{aligned}$$

(each line has vectors of equal norm). The next step is to find the set of irreducible vectors $\text{Irr}(\mathcal{L})$. We consider the subset of $R(\mathcal{L})$ containing relevant vectors which cannot be written as a sum of two strictly shorter vectors (by cross-checking

with the set of relevant vectors). It turns out that this set just contains all the vectors achieving the successive minima thus it must be that this is $\text{Irr}(\mathcal{L})$.

$$\begin{aligned} \text{Irr}(\mathcal{L}) = \{ & \pm (0, -1, 0, 1, 0), \\ & \pm (0, -1, 0, -1, 1), \\ & \pm (0, 0, 2, 0, 0), \\ & \pm (-1, 0, -1, 1, 2), \pm(-1, 0, 1, 1, 2), \pm(-1, 1, -1, 0, 2), \pm(-1, 1, 1, 0, 2), \\ & \pm (3, 1, 0, 0, 1), \pm(3, 0, 0, 1, 1)\} \end{aligned}$$

The set $\text{Irr}(\mathcal{L})$ contains 5 equivalence classes according to the equivalence relation (13). We denote them by C_i for $i = 1, \dots, 5$. As we can see for the first three of them, computing a set of representatives $\tilde{C}_1, \tilde{C}_2, \tilde{C}_3$ is trivial as in these cases it will be $\tilde{C}_1 = C_1, \tilde{C}_2 = C_2$ and $\tilde{C}_3 = C_3$. The cases of C_4 and C_5 are more interesting. We start by examining C_5 as it contains less vectors. We set $\mathbf{v}_1 = (3, 1, 0, 0, 1)$ and $\mathbf{v}_2 = (3, 0, 0, 1, 1)$. Next we draw the corresponding graph with vertices the $\pm\mathbf{v}_1, \pm\mathbf{v}_2$ and edges if the pairwise differences are strictly shorter.




Fig. 1: The graph of the class C_5

The graph in Figure 1 shows that we can take either $\tilde{C}_5 = \{\pm\mathbf{v}_1\}$ or $\tilde{C}_5 = \{\pm\mathbf{v}_2\}$. We are now going to do the same for the class C_4 . We set $\mathbf{v}_1 = (-1, 0, -1, 1, 2), \mathbf{v}_2 = (-1, 0, 1, 1, 2), \mathbf{v}_3 = (-1, 1, -1, 0, 2), \mathbf{v}_4 = (-1, 1, 1, 0, 2)$.




Fig. 2: The graph of the class C_4

The graph in Figure 2 shows that we can take $\tilde{C}_4 = \{\pm \mathbf{v}_i\}$ for any $i = 1, 2, 3, 4$. Therefore one choice for the set $P(\mathcal{L})$ is the following.

$$P(\mathcal{L}) = \{ \pm (0, -1, 0, 1, 0), \\ \pm (0, -1, 0, -1, 1), \\ \pm (0, 0, 2, 0, 0), \\ \pm (-1, 0, -1, 1, 2), \\ \pm (3, 1, 0, 0, 1) \}$$

Remark 12. The above example should not mislead the reader to the assumption that the corresponding graph of each equivalence class will always have at least two connected components. It can happen that the graph of a class is connected. One such example can be derived from the family of lattices examined in Theorem 10.

One property of the set $\text{Irr}(\mathcal{L})$ was that it includes a generating set of \mathcal{L} . We can show that $P(\mathcal{L})$ inherits that property.

Proposition 7 (Complete system generates). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . Then for every $P(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L})$ there exists a generating set \mathbf{G} of \mathcal{L} such that $\mathbf{G} \subseteq P(\mathcal{L})$.*

Proof. As in the proof of Proposition 5 for the set $\text{Irr}(\mathcal{L})$ we will prove that the set $P(\mathcal{L})$ spans the lattice and therefore it includes a generating set. However, in this case the proof is more technical. Let $P(\mathcal{L})$ be a complete system of irreducible vectors of \mathcal{L} as defined in (14). We have already shown that $\text{Irr}(\mathcal{L})$ is finite as $\text{Irr}(\mathcal{L}) \subseteq R(\mathcal{L})$ and thus we can define $t := |\text{Irr}(\mathcal{L})/\sim|$. We further set C_i for $i = 1, \dots, t$ to be the equivalence classes in $\text{Irr}(\mathcal{L})/\sim$. Hence, the set $P(\mathcal{L})$ can be written as $P(\mathcal{L}) = \cup_{i=1}^t \tilde{C}_i$. Each equivalence class C_i contains all irreducible vectors of a specific length μ_i , and we can assume that we have ordered the C_i according to increasing μ_i . We define the following sequence of subsets of $\text{Irr}(\mathcal{L})$:

$$A_i := \left(\bigcup_{j=1}^{i-1} C_j \right) \cup \left(\bigcup_{j=i}^t \tilde{C}_j \right) \quad \text{for } i = 1, \dots, t+1.$$

As for every i it holds $\tilde{C}_i \subseteq C_i$ then it follows that $A_i(\mathcal{L}) \subseteq A_{i+1}(\mathcal{L})$ and thus

$$P(\mathcal{L}) = A_1 \subseteq A_2 \subseteq \dots \subseteq A_t \subseteq A_{t+1} = \text{Irr}(\mathcal{L}).$$

We will prove by induction that each term of this sequence of sets spans the lattice \mathcal{L} .

Base case $i = t+1$: The set $\text{Irr}(\mathcal{L}) = A_{t+1}$ spans the lattice as it was already shown in Proposition 5.

Induction hypothesis: Assume that it holds for some $i = k$, i.e. A_k spans the lattice for some $k \in \{2, \dots, t+1\}$.

Induction step: Prove that A_{k-1} spans the lattice. By the definition of the sets

A_i we can conclude that $A_{k-1} = A_k \setminus (C_{k-1} \setminus \tilde{C}_{k-1})$. By the induction hypothesis it suffices to show that the vectors in $C_{k-1} \setminus \tilde{C}_{k-1}$ can be generated by the vectors in A_{k-1} . Let $\mathbf{v} \in C_{k-1} \setminus \tilde{C}_{k-1}$. As $\mathbf{v} \in C_{k-1}$ but $\mathbf{v} \notin \tilde{C}_{k-1}$ this implies that there exists a $\tilde{\mathbf{v}} \in \tilde{C}_{k-1}$ such that $\|\mathbf{v} + \tilde{\mathbf{v}}\| < \|\mathbf{v}\|$. This holds because \tilde{C}_{k-1} is maximal by definition. We set $\mathbf{w} = \mathbf{v} + \tilde{\mathbf{v}}$. Furthermore as $\|\mathbf{w}\| < \|\mathbf{v}\|$ then \mathbf{w} is either irreducible or can be written as a sum of irreducible vectors shorter than $\|\mathbf{v}\|$. We use the ordering of the C_i . Thus by its definition the set A_{k-1} contains all the vectors in $\text{Irr}(\mathcal{L})$ which are shorter than $\|\mathbf{v}\|$. Hence as, $\|\mathbf{w}\| < \|\mathbf{v}\|$ this implies that \mathbf{w} can be generated by the vectors in A_{k-1} . So, concluding we wrote \mathbf{v} as $\mathbf{v} = \mathbf{w} - \tilde{\mathbf{v}}$ where both \mathbf{w} and $\tilde{\mathbf{v}}$ belong to A_{k-1} . This concludes the proof. \square

For $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{L}$ we denote by $\vartheta(\mathbf{v}_1, \mathbf{v}_2)$ the angle formed by $\mathbf{v}_1, \mathbf{v}_2$.

Proposition 8 (Properties of complete system). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n , and $\mathbf{p}_1, \mathbf{p}_2 \in P(\mathcal{L})$ such that $\mathbf{p}_1 \neq \pm\mathbf{p}_2$. Then it holds that*

- (i) $\min\{\|\mathbf{p}_1 \pm \mathbf{p}_2\|\} \geq \max\{\|\mathbf{p}_1\|, \|\mathbf{p}_2\|\}$ and
- (ii) $|\cos \vartheta(\mathbf{p}_1, \mathbf{p}_2)| \leq \frac{1}{2}$.

Proof. (Part i) By Lemma 4 we have that

$$\text{Irr}(\mathcal{L}) = \{\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\} \mid \forall \mathbf{x} \in \mathcal{L} \text{ with } \|\mathbf{x}\| < \|\mathbf{v}\| \text{ it holds } \|\mathbf{v} - \mathbf{x}\| \geq \|\mathbf{v}\|\}.$$

Let $\mathbf{p}_1, \mathbf{p}_2 \in P(\mathcal{L})$ such that $\mathbf{p}_1 \neq \pm\mathbf{p}_2$. Without loss of generality we assume that $\|\mathbf{p}_2\| \leq \|\mathbf{p}_1\|$. Initially we will prove that $\|\mathbf{p}_1 + \mathbf{p}_2\| \geq \max\{\|\mathbf{p}_1\|, \|\mathbf{p}_2\|\}$.

Case 1: If $\|\mathbf{p}_2\| < \|\mathbf{p}_1\|$. Then $\mathbf{p}_1, \mathbf{p}_2 \in \text{Irr}(\mathcal{L})$ and they are not in the same class. Using the description of the lemma with $\mathbf{v} = \mathbf{p}_1 \in \text{Irr}(\mathcal{L})$ and $\mathbf{x} = -\mathbf{p}_2$ we get $\|\mathbf{p}_1 + \mathbf{p}_2\| \geq \|\mathbf{p}_1\|$. But as $\|\mathbf{p}_2\| < \|\mathbf{p}_1\|$ we can conclude that $\|\mathbf{p}_1 + \mathbf{p}_2\| \geq \max\{\|\mathbf{p}_1\|, \|\mathbf{p}_2\|\}$.

Case 2: If $\|\mathbf{p}_2\| = \|\mathbf{p}_1\|$. Then $\mathbf{p}_1, \mathbf{p}_2 \in \text{Irr}(\mathcal{L})$ and they are in the same class. Let $S \in \text{Irr}(\mathcal{L})/\sim$ such that $\mathbf{p}_1, \mathbf{p}_2 \in S$. Then as $\mathbf{p}_1, \mathbf{p}_2 \in P(\mathcal{L})$ we get that $\mathbf{p}_1, \mathbf{p}_2$ belong to the same \tilde{S} . Thus, by the definition of \tilde{S} we can again conclude that $\|\mathbf{p}_1 + \mathbf{p}_2\| \geq \max\{\|\mathbf{p}_1\|, \|\mathbf{p}_2\|\}$.

The result follows from the fact that for every $\mathbf{v} \in P(\mathcal{L})$ also $-\mathbf{v} \in P(\mathcal{L})$.

(Part ii) Let $\mathbf{p}_1, \mathbf{p}_2 \in P(\mathcal{L})$ such that $\mathbf{p}_1 \neq \pm\mathbf{p}_2$. Without loss of generality we assume that $\|\mathbf{p}_2\| \leq \|\mathbf{p}_1\|$. By part (i) we get that $\|\mathbf{p}_1 \pm \mathbf{p}_2\| \geq \|\mathbf{p}_1\|$. This in turn implies that $2|\langle \mathbf{p}_1, \mathbf{p}_2 \rangle| \leq \|\mathbf{p}_2\|^2$. Hence,

$$|\cos \vartheta(\mathbf{p}_1, \mathbf{p}_2)| = \frac{|\langle \mathbf{p}_1, \mathbf{p}_2 \rangle|}{\|\mathbf{p}_1\| \|\mathbf{p}_2\|} \leq \frac{|\langle \mathbf{p}_1, \mathbf{p}_2 \rangle|}{\|\mathbf{p}_2\|^2} \leq \frac{1}{2}.$$

\square

We will use the following theorem in order to bound $|P(\mathcal{L})|$.

Theorem 7 (Upper bound on kissing constant [17]). *Let $A(n, \phi_0)$ be the maximal size of any set C of points in \mathbb{R}^n such that the angle between any two*

distinct vectors $\mathbf{v}_i, \mathbf{v}_j \in C$ (denoted $\phi_{\mathbf{v}_i, \mathbf{v}_j}$) is at least ϕ_0 . If $0 < \phi_0 < 63^\circ$, then for all sufficiently large n , $A(n, \phi_0) = 2^{cn}$ for some

$$c \leq -\frac{1}{2} \log_2(1 - \cos(\phi_0)) - 0.099. \quad (15)$$

Proposition 9 (Upper bound on $|\mathsf{P}(\mathcal{L})|$). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . It holds that $|\mathsf{P}(\mathcal{L})| < 2^{0.402n}$.*

Proof. By using Theorem 7 with $\phi_0 = \frac{\pi}{3}$ (which can be deduced from Proposition 8) we get that $|\mathsf{P}(\mathcal{L})| = 2^{cn}$ with $c \leq -\frac{1}{2} \log_2(1 - \cos(\frac{\pi}{3})) - 0.099$. Evaluating the right hand side of this inequality implies the result. \square

Proposition 8 states the same condition that is also satisfied by the output of the GaussSieve algorithm described in [23]. As in the paper describing the GaussSieve algorithm [23] the size of $\mathsf{P}(\mathcal{L})$ can actually be bounded by the kissing number τ_n . Following the same arguments as in [23] we can argue that in practice we expect $|\mathsf{P}(\mathcal{L})| \approx 2^{0.21n}$ which is a factor 2 smaller in the exponent than the provable bound $|\mathsf{P}(\mathcal{L})| < 2^{0.402n}$.

A result that might be of interest in the search for lattices reaching the kissing number is the following.

Theorem 8 (Lattices achieving the kissing number). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . If the lattice \mathcal{L} is such that it reaches the kissing number τ_n then $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L})$.*

Proof. As the lattice \mathcal{L} reaches the kissing number τ_n , that implies $|S_1(\mathcal{L})| = \tau_n$. By Proposition 8 we can conclude that the angle between any two vectors in $\mathsf{P}(\mathcal{L})$ is at least $\pi/3$. This is also the minimal possible angle between the centers of two equal n -dimensional spheres which touch another central sphere of the same size without intersecting. Hence $|\mathsf{P}(\mathcal{L})| \leq \tau_n$. Combining this with $S_1(\mathcal{L}) \subseteq \mathsf{P}(\mathcal{L})$ and $|S_1(\mathcal{L})| = \tau_n$ implies that $\mathsf{P}(\mathcal{L}) = S_1(\mathcal{L})$. As the set $\mathsf{P}(\mathcal{L})$ was build from classes of $\text{Irr}(\mathcal{L})$ and we showed that it actually contains only vectors of norm $\lambda_1(\mathcal{L})$ that means that there is only one class in $\text{Irr}(\mathcal{L})/\sim$, namely the class of $S_1(\mathcal{L})$. But in this class there is no pair of vectors that adds to a shorter one, thus the whole class is included in $\mathsf{P}(\mathcal{L})$. That implies that $\text{Irr}(\mathcal{L}) = \mathsf{P}(\mathcal{L}) = S_1(\mathcal{L})$. \square

Remark 13. A similar result for the set $\mathsf{R}(\mathcal{L})$ is not possible. For example for the root lattice E_8 reaching the kissing number in dimension 8 it holds $S_1(E_8) = \mathsf{R}(E_8)$ but for the Leech lattice A_{24} it holds that $S_1(A_{24}) \subsetneq \mathsf{R}(A_{24})$ (see [8]).

5 Computation of the set $\mathsf{P}(\mathcal{L})$

In the previous sections we investigated some properties of the set $\mathsf{P}(\mathcal{L})$ and its relation to the set $\mathsf{R}(\mathcal{L})$. Ultimately we aim in using this set instead of $\mathsf{R}(\mathcal{L})$ due to its provably smaller cardinality. However, in order to actually benefit from this replacement an algorithm that computes $\mathsf{P}(\mathcal{L})$ without using the set $\mathsf{R}(\mathcal{L})$ is needed. The goal of this section is to examine ways of computing the set $\mathsf{P}(\mathcal{L})$.

5.1 The “brute force” approach

If the set $\text{Irr}(\mathcal{L})$ is given then the set $\text{P}(\mathcal{L})$ can be computed by means of a graph-based technique already described in Remark 9 and further analysed in Appendix B. Thus, it suffices to describe an algorithm which computes the set $\text{Irr}(\mathcal{L})$. The naive approach is to use the fact that $\text{Irr}(\mathcal{L}) \subseteq \text{R}(\mathcal{L})$. Hence, as a first step one can run the algorithm described in [22] in order to get the set $\text{R}(\mathcal{L})$. Then having a superset of $\text{Irr}(\mathcal{L})$ it suffices to remove all the reducible vectors from it. This can be done by iterating through $\text{R}(\mathcal{L})$ and checking for each $\mathbf{r} \in \text{R}(\mathcal{L})$ if there exists a $\mathbf{v} \in \text{R}(\mathcal{L})$ such that $\|\mathbf{v}\| < \|\mathbf{r}\|$ and $\|\mathbf{r} - \mathbf{v}\| < \|\mathbf{r}\|$. If $\mathbf{r} \in \text{Irr}(\mathcal{L})$ then by definition there will not exist a vector $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\| < \|\mathbf{r}\|$ and $\|\mathbf{r} - \mathbf{v}\| < \|\mathbf{r}\|$ and thus the algorithm will not discard any of the irreducible vectors. If the vector \mathbf{r} is reducible then we need the following heuristic assumption.

Heuristic Assumption 1 (Witness of reducibility) *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n with $\text{Irr}(\mathcal{L}) \neq \text{R}(\mathcal{L})$. If $\mathbf{r} \in \text{R}(\mathcal{L}) \setminus \text{Irr}(\mathcal{L})$ then $\exists \mathbf{v} \in \text{R}(\mathcal{L})$ such that $\|\mathbf{v}\| < \|\mathbf{r}\|$ and $\|\mathbf{r} - \mathbf{v}\| < \|\mathbf{r}\|$.*

Heuristic assumption 1 can be considered as the analogue of Lemma 1 for the set $\text{Irr}(\mathcal{L})$. Lemma 1 guaranteed that for every non-relevant vector there would exist a relevant vector acting as a “witness” of “non-relevancy”. Heuristic assumption 1 speculates that for every reducible relevant vector there exists a relevant vector acting as a “witness” of reducibility.

5.2 Using the GaussSieve/MinkowskiSieve algorithms

As it was already mentioned in Section 4.2, it is expected that the output of the GaussSieve algorithm [23] will be closely related to a set $\text{P}(\mathcal{L})$. This conjecture was motivated by the fact that both sets, $\text{P}(\mathcal{L})$ and the output of the GaussSieve, possess the property $\min\{\|\mathbf{v}_1 \pm \mathbf{v}_2\|\} \geq \max\{\|\mathbf{v}_1\|, \|\mathbf{v}_2\|\}$ for any pair of $\mathbf{v}_1 \neq \pm \mathbf{v}_2$ in the set. At this point it should be clarified that for our purposes we will consider a slightly modified version of the GaussSieve algorithm which will be described here.

For our purposes we will use the GaussSieve algorithm 1 but with the modified version of the GaussReduce function 2. In this way the following conditions are met.

- (i) Any irreducible vector which has been added to the GaussSieve list L will never be removed from it.
- (ii) Any irreducible vector encountered by the algorithm will be added to L provided that it can extend its class representative set already in L .

Lemma 5 (Modified GaussSieve algorithm). *The GaussSieve algorithm 1 equipped with the function PrimeGaussReduce (Algorithm 2) satisfies both properties (i) and (ii).*

¹ This version is the one used in [32] as well.

Algorithm 1 The GaussSieve algorithm as described in [23]

Require: A basis \mathbf{B} of a lattice $\mathcal{L}(\mathbf{B})$ and a $c > 0$.**Ensure:** A list $L \subset \mathcal{L}$ s.t. $\min\{\|\mathbf{v}_1 \pm \mathbf{v}_2\|\} \geq \max\{\|\mathbf{v}_1\|, \|\mathbf{v}_2\|\}$ for all $\mathbf{v}_1, \mathbf{v}_2 \in L$.

<pre> function GAUSSSIEVE(\mathbf{B}, c) $L \leftarrow \{\mathbf{0}\}, S \leftarrow \{\}, K \leftarrow 0$ while $K < c$ do if S is not empty then $\mathbf{v}_{\text{new}} \leftarrow S.\text{pop}()$ else $\mathbf{v}_{\text{new}} \leftarrow \text{SampleGaussian}(\mathbf{B})$ end if $\mathbf{v}_{\text{new}} \leftarrow \text{GaussReduce}(\mathbf{v}_{\text{new}}, L, S)$ if $\mathbf{v}_{\text{new}} = \mathbf{0}$ then $K \leftarrow K + 1$ else $L \leftarrow L \cup \{\mathbf{v}_{\text{new}}\}$ end if end while end function </pre>	<pre> function GAUSSREDUCE(\mathbf{p}, L, S) while $\exists \mathbf{v}_i \in L : \ \mathbf{v}_i\ \leq \ \mathbf{p}\$ and $\ \mathbf{p} - \mathbf{v}_i\ \leq \ \mathbf{p}\$ do $\mathbf{p} \leftarrow \mathbf{p} - \mathbf{v}_i$ end while while $\exists \mathbf{v}_i \in L : \ \mathbf{v}_i\ > \ \mathbf{p}\$ and $\ \mathbf{v}_i - \mathbf{p}\ \leq \ \mathbf{v}_i\$ do $L \leftarrow L \setminus \{\mathbf{v}_i\}$ $S.\text{push}(\mathbf{v}_i - \mathbf{p})$ end while return \mathbf{p} end function </pre>
---	---

Proof. (Property *i*) The only way for a vector $\mathbf{v}_i \in L$ to be removed from the list L is by entering the **while** loop in line 5 of the PrimeGaussReduce function. Let $\mathbf{v}_i \in L$ and also $\mathbf{v}_i \in \text{Irr}(\mathcal{L})$. In order for the algorithm to remove \mathbf{v}_i from L it should encounter another vector \mathbf{p} such that $\|\mathbf{v}_i\| > \|\mathbf{p}\|$ and $\|\mathbf{v}_i - \mathbf{p}\| < \|\mathbf{v}_i\|$ or $\|\mathbf{v}_i\| > \|\mathbf{p}\|$ and $\|\mathbf{v}_i + \mathbf{p}\| < \|\mathbf{v}_i\|$ which contradicts the irreducibility of \mathbf{v}_i . (Property *ii*) Assume that the function PrimeGaussReduce is called and in some iteration of the **while** loop in line 2, \mathbf{p} becomes such that $\mathbf{p} \in \text{Irr}(\mathcal{L})$. In order for \mathbf{p} to not be added in L this would mean that \mathbf{p} could be further modified by the **while** loop in line 2. Thus the algorithm should encounter another vector $\mathbf{v}_i \in L$ such that $\|\mathbf{v}_i\| \leq \|\mathbf{p}\|$ and $\|\mathbf{p} \pm \mathbf{v}_i\| < \|\mathbf{p}\|$. The case where $\|\mathbf{v}_i\| < \|\mathbf{p}\|$ and $\|\mathbf{p} \pm \mathbf{v}_i\| < \|\mathbf{p}\|$ violates the irreducibility of \mathbf{p} and thus can be disregarded. This leaves only one possible case, namely $\|\mathbf{v}_i\| = \|\mathbf{p}\|$ and $\|\mathbf{p} \pm \mathbf{v}_i\| < \|\mathbf{p}\|$. This condition implies that \mathbf{v}_i and \mathbf{p} belong to the same equivalence class and they are adjacent. Therefore this pair of vectors cannot belong to any set $P(\mathcal{L})$ of \mathcal{L} . Hence \mathbf{p} should not be included in L anyway and the algorithm correctly further reduces it. \square

Remark 14. If the PrimeGaussReduce function in line 5 was the same as in the original GaussReduce, then the algorithm could encounter an instance where it would enter the loop with $\|\mathbf{v}_i\| > \|\mathbf{p}\|$, $\|\mathbf{v}_i - \mathbf{p}\| = \|\mathbf{v}_i\|$ and $\mathbf{v}_i, \mathbf{v}_i - \mathbf{p} \in \text{Irr}(\mathcal{L})$. This could be possible if an equivalence class in $\text{Irr}(\mathcal{L})$ was not trivial. In this case the algorithm would remove the vector \mathbf{v}_i from the list and add its equivalent $\mathbf{v}_i - \mathbf{p}$ to S . As a result for these non-trivial classes the algorithm could behave

Algorithm 2 The modified GaussReduce function

```

1: function PRIMEGAUSSREDUCE( $\mathbf{p}, L, S$ )1
2:   while  $\exists \mathbf{v}_i \in L : \|\mathbf{v}_i\| \leq \|\mathbf{p}\| \wedge \|\mathbf{p} \pm \mathbf{v}_i\| < \|\mathbf{p}\|$  do
3:      $\mathbf{p} \leftarrow \mathbf{p} \pm \mathbf{v}_i$ 
4:   end while
5:   while  $\exists \mathbf{v}_i \in L : \|\mathbf{v}_i\| > \|\mathbf{p}\| \wedge \|\mathbf{v}_i \pm \mathbf{p}\| < \|\mathbf{v}_i\|$  do
6:      $L \leftarrow L \setminus \{\mathbf{v}_i\}$ 
7:     S.push( $\mathbf{v}_i \pm \mathbf{p}$ )
8:   end while
9:   return  $\mathbf{p}$ 
10: end function

```

in a bad way by repetitively removing and adding representatives of the same class.

Remark 15. If the PrimeGaussReduce function in line 2 was the same as in the original GaussReduce, then the algorithm could encounter an instance where it would enter the loop with $\|\mathbf{v}_i\| \leq \|\mathbf{p}\|$, $\|\mathbf{p} - \mathbf{v}_i\| = \|\mathbf{p}\|$ and $\mathbf{p}, \mathbf{p} - \mathbf{v}_i \in \text{Irr}(\mathcal{L})$. Thus, \mathbf{p} and $\mathbf{p} - \mathbf{v}_i$ are equivalent. In case $\|\mathbf{v}_i\| < \|\mathbf{p}\|$ then \mathbf{p} and $\mathbf{p} - \mathbf{v}_i$ are also adjacent in the class graph and therefore in this case the algorithm would cycle through the adjacent vectors of \mathbf{p} . Therefore there is no need to perform a reduction in this case. In case $\|\mathbf{v}_i\| = \|\mathbf{p}\|$ then all three $\mathbf{p}, \mathbf{v}_i, \mathbf{p} - \mathbf{v}_i$ are equivalent but not adjacent. Hence in this case the algorithm does not make any progress by replacing \mathbf{p} by $\mathbf{p} - \mathbf{v}_i$. Thus, there is no need to perform a reduction in this case as well.²

We consider the GaussSieve algorithm 1 equipped with the PrimeGaussReduce function (Algorithm 2). We denote by GaussSieve(\mathcal{L}) a list of vectors L created by this algorithm and possessing the property that L cannot be further modified by the algorithm. In order to relate the sets GaussSieve(\mathcal{L}) and $P(\mathcal{L})$ we give the following definition.

Definition 10 (Partitioning $P(\mathcal{L})$ by sign). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . Given a $P(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L})$ we define $P^+(\mathcal{L})$ and $P^-(\mathcal{L})$ to be a partition of $P(\mathcal{L})$ according to sign.*

In other words, we take for $P^+(\mathcal{L})$ some subset of $P(\mathcal{L})$ such that of each pair $\pm \mathbf{v} \in P^+(\mathcal{L})$ exactly one is in $P^+(\mathcal{L})$. Of course, there are many choices for $P^+(\mathcal{L})$ and $P^-(\mathcal{L})$, any one will do.

Even though the GaussSieve algorithm converges to a set which is maximal in \mathcal{L} under the property $\min\{\|\mathbf{v}_1 \pm \mathbf{v}_2\|\} \geq \max\{\|\mathbf{v}_1\|, \|\mathbf{v}_2\|\}$, the same is not true in general for the set $P^+(\mathcal{L})$ as shown by experiments (Table 1). In particular, we can conclude by Lemma 5 that if we allow this modified version of the GaussSieve

² However, as $\mathbf{p} - \mathbf{v}_i$ is not adjacent to both \mathbf{p} and \mathbf{v}_i an option could be to move $\mathbf{p} - \mathbf{v}_i$ to the stack S for further consideration later.

to run long enough i.e. it samples “enough” vectors, then the output will converge to a set $\text{GaussSieve}(\mathcal{L})$, which will contain a $P^+(\mathcal{L})$.

Hence we cannot claim that the GaussSieve algorithm converges to a set $P^+(\mathcal{L})$ but only to a superset of it. The fact that a $P^+(\mathcal{L})$ is not maximal in \mathcal{L} under the property $\min\{\|\mathbf{v}_1 \pm \mathbf{v}_2\|\} \geq \max\{\|\mathbf{v}_1\|, \|\mathbf{v}_2\|\}$ implies the existence of vectors which are not irreducible but they also cannot be reduced by any of the vectors in $P(\mathcal{L})$.

The definition of the set $P_2(\mathcal{L})$ will help us in bounding the output of the GaussSieve algorithm. Also, the definition of the sets $P_k(\mathcal{L})$ for $k > 2$ will help us in bounding the output of modified versions of “higher” sieving algorithms like the Triple and Quadruple MinkowskiSieve, described in [3].

Definition 11 (Pairwise irreducible system). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . Given a $P(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L})$ we define*

$$P_2(\mathcal{L}) := \{\mathbf{v} \in \mathcal{L} \mid \nexists \mathbf{p} \in P(\mathcal{L}) \text{ with } \|\mathbf{p}\| < \|\mathbf{v}\| \text{ and } \|\mathbf{v} - \mathbf{p}\| < \|\mathbf{v}\|\}.$$

A first remark on this definition is that as $P(\mathcal{L}) \subseteq \mathcal{L}$ then $P(\mathcal{L}) \subseteq P_2(\mathcal{L})$. The (modified) GaussSieve algorithm converges to a set $\text{GaussSieve}(\mathcal{L})$ including a set $P^+(\mathcal{L})$ hence, $\text{GaussSieve}(\mathcal{L})$ can be bounded as follows.

$$P^+(\mathcal{L}) \subseteq \text{GaussSieve}(\mathcal{L}) \subseteq P_2(\mathcal{L}) \quad (16)$$

Under this set inequality $\text{GaussSieve}(\mathcal{L})$ can be viewed in the following way. A set $\text{GaussSieve}(\mathcal{L})$ can be considered as the closure of a $P^+(\mathcal{L})$ in $P_2(\mathcal{L})$ under the property of Gauss-reduction. In more detail $\text{GaussSieve}(\mathcal{L})$ can be viewed as the minimal (according to included vector norms) subset of a $P_2(\mathcal{L})$ including $P^+(\mathcal{L})$ and being a maximal subset of $P_2(\mathcal{L})$ with the property of Gauss-reduction (i.e. $\min\{\|\mathbf{v}_1 \pm \mathbf{v}_2\|\} \geq \max\{\|\mathbf{v}_1\|, \|\mathbf{v}_2\|\}$).

Definition 12 (k -wise irreducible system). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $k \in \mathbb{N}$ with $k \geq 2$. Given a $P(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L})$ we define*

$$P_{k+1}(\mathcal{L}) := \{\mathbf{v} \in P_k(\mathcal{L}) \mid \nexists \mathbf{p} \in P^{(k)}(\mathcal{L}) \text{ with } \|\mathbf{p}\| < \|\mathbf{v}\| \text{ and } \|\mathbf{v} - \mathbf{p}\| < \|\mathbf{v}\|\}$$

where $P^{(k)}(\mathcal{L})$ is defined as

$$\bigcup_{i=1}^{\lfloor k/2 \rfloor} \{\mathbf{v}_1 + \mathbf{v}_2 \mid \mathbf{v}_1 \in P^{(i)}(\mathcal{L}), \mathbf{v}_2 \in P^{(k-i)}(\mathcal{L}) \text{ and } \|\mathbf{v}_j\| < \|\mathbf{v}_1 + \mathbf{v}_2\|, \\ \|\mathbf{v}_l\| \leq \|\mathbf{v}_1 + \mathbf{v}_2\| \text{ where } j, l \in \{1, 2\}\}$$

and $P^{(1)}(\mathcal{L}) := P(\mathcal{L})$.

Lemma 6 (Relating k -wise irreducible systems). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $P(\mathcal{L})$ be a subset of $\text{Irr}(\mathcal{L})$. Then for the sequence $P_k(\mathcal{L})$ given in definition 12 it holds that*

- (i) $P_{k+1}(\mathcal{L}) \subseteq P_k(\mathcal{L})$ for every $k \geq 2$.
- (ii) $\lim_{k \rightarrow \infty} P_k(\mathcal{L}) = \text{Irr}(\mathcal{L})$.

So, in one line: there exists an $l \geq 2$ such that

$$P_2(\mathcal{L}) \supseteq \dots \supseteq P_k(\mathcal{L}) \supseteq P_{k+1}(\mathcal{L}) \supseteq \dots \supseteq P_l(\mathcal{L}) = P_{l+1}(\mathcal{L}) = \dots = \text{Irr}(\mathcal{L}).$$

Proof. Part (i) of the lemma is an immediate consequence of the $P_k(\mathcal{L})$ definition. Initially we show that $\text{Irr}(\mathcal{L}) \subseteq P_k(\mathcal{L})$ for every $k \geq 2$. This follows directly by the definition of $P_k(\mathcal{L})$ and the fact that $P^{(k)}(\mathcal{L}) \subseteq \mathcal{L}$. By the (recursive) definition of $P_k(\mathcal{L})$ it follows that it includes all vectors $\mathbf{v} \in \mathcal{L}$ such that they can not be reduced by any shorter vector in $\cup_{i=1}^{k-1} P^{(i)}(\mathcal{L})$. Thus for part (iii) of the lemma it suffices to show that $\lim_{k \rightarrow \infty} \cup_{i=1}^k P^{(i)}(\mathcal{L}) = \mathcal{L}$. As $P^{(i)}(\mathcal{L}) \subseteq \mathcal{L}$ for every $i \geq 1$ it follows that $\lim_{k \rightarrow \infty} \cup_{i=1}^k P^{(i)}(\mathcal{L}) \subseteq \mathcal{L}$. It is only left proving the converse inequality. Let $\mathbf{v} \in \mathcal{L}$, it suffices to show that $\exists k \geq 1$ such that $\mathbf{v} \in P^{(k)}(\mathcal{L})$.

A vector $\mathbf{v} \in \mathcal{L}$ can be repeatedly reduced as in the proof of Proposition 5 until it is written as a sum $\mathbf{v} = \sum_{i=1}^l \mathbf{p}_i$ of shorter vectors $\mathbf{p}_i \in \text{Irr}(\mathcal{L})$ for some $l \geq 1$. This decomposition satisfies the recursive condition implied by the definition of the $P^{(k)}(\mathcal{L})$. If all the vectors $\mathbf{p}_i \in \text{Irr}(\mathcal{L})$ actually belong to $P(\mathcal{L})$ then $\mathbf{v} \in P^{(l)}(\mathcal{L})$ and we are done. If there exists some $\mathbf{p}_i \in \text{Irr}(\mathcal{L}) \setminus P(\mathcal{L})$ then $\mathbf{p}_i = \tilde{\mathbf{p}}_i + \mathbf{p}'_i$ where $\tilde{\mathbf{p}}_i \in P(\mathcal{L})$ and $\|\mathbf{p}'_i\| < \|\mathbf{p}_i\|$, $\|\mathbf{p}_i\| = \|\tilde{\mathbf{p}}_i\|$ by the definition of $P(\mathcal{L})$. Thus, \mathbf{p}'_i can be further get decomposed in shorter vectors (like \mathbf{v}) and as $\|\mathbf{p}'_i\| < \|\mathbf{p}_i\|$ progress was made which implies that this decomposition will finish after finitely many steps. Therefore \mathbf{v} can be repeatedly reduced until it is written as a sum of vectors in $P(\mathcal{L})$, concluding the proof. \square

We are now going to describe the “higher” sieving algorithms which we will consider. We have already mentioned the Triple and the Quadruple MinkowskiSieve described in [3]. The difference between the GaussSieve algorithm and these higher ones lies in the reduction function. Hence, if we equip Algorithm 1 with function PrimeMinkowskiReduce (Algorithm 3), we get the modified MinkowskiSieve which we are interested in.

The modification compared to the description in [3] appears in lines 10 and 21 of Algorithm 3, where the extra conditions $\|\mathbf{w}\| \leq \|\mathbf{p}\|$ and $\|\mathbf{w}\| < \|\mathbf{v}_{k-1}\|$ respectively are added. By adding these conditions it is guaranteed to get an output list which will satisfy properties (i) and (ii) like in Lemma 5 for the GaussSieve. Hence, based on these properties it can be concluded that the output list of vectors will again contain a set $P^+(\mathcal{L})$. In order to ease our exposition we set the following notation.

Let $k \in \mathbb{N}$ with $k \geq 2$. We consider the k -MinkowskiSieve algorithm equipped with the function PrimeMinkowskiReduce (Algorithm 3). We denote by $\text{MinkowskiSieve}_k(\mathcal{L})$ a list of vectors L created by this algorithm and possessing the property that L cannot be further modified by the algorithm. Note that for $k = 2$ one has $\text{MinkowskiSieve}_2(\mathcal{L}) = \text{GaussSieve}(\mathcal{L})$.

³ We assume $\|\mathbf{v}_i\| \leq \|\mathbf{v}_{i+1}\|$.

Algorithm 3 The modified MinkowskiReduce function

```

1: function PRIMEMINKOWSKIREDUCE( $\mathbf{p}, L, S, k$ )
2:   loop = true
3:   while loop do
4:     loop = false
5:     if  $k > 2$  then
6:       PRIMEMINKOWSKIREDUCE( $\mathbf{p}, L, S, k - 1$ )
7:     end if
8:     for all  $\{\mathbf{v}_1, \dots, \mathbf{v}_{k-1}\} \subset L$  s.t.  $\|\mathbf{v}_i\| \leq \|\mathbf{p}\|$  do
9:       for all  $\mathbf{w} \in \{\sum_{i=1}^{k-1} (-1)^{a_i} \mathbf{v}_i\}$  do
10:        if  $\|\mathbf{w}\| \leq \|\mathbf{p}\|$  and  $\|\mathbf{p} - \mathbf{w}\| < \|\mathbf{p}\|$  then
11:           $\mathbf{p} \leftarrow \mathbf{p} - \mathbf{w}$ 
12:          loop = true
13:          goto next
14:        end if
15:      end for
16:    end for
17:    next:
18:  end while
19:  for all  $\{\mathbf{v}_1, \dots, \mathbf{v}_{k-1}\} \subset L$  s.t.  $\|\mathbf{v}_{k-1}\| > \|\mathbf{p}\|$  do
20:    for all  $\mathbf{w} \in \{(-1)^{a_0} \mathbf{p} + \sum_{i=1}^{k-2} (-1)^{a_i} \mathbf{v}_i\}$  do
21:      if  $\|\mathbf{w}\| < \|\mathbf{v}_{k-1}\|$  and  $\|\mathbf{v}_{k-1} - \mathbf{w}\| < \|\mathbf{v}_{k-1}\|$  then
22:         $L \leftarrow L \setminus \{\mathbf{v}_{k-1}\}$ 
23:        S.push( $\mathbf{v}_{k-1} - \mathbf{w}$ )
24:      end if
25:    end for
26:  end for
27:  return  $\mathbf{p}$ 
28: end function

```

Remark 16. The output of the modified k -MinkowskiSieve algorithm will not be a list of vectors which will be k -Minkowski-reduced if $k > 2$ (for the Minkowski-reduced definition see [26]). If this was desired, then the lines 10 and 21 of Algorithm 3 should be modified in order to allow reductions by longer vectors as well. For a k -Minkowski-reduced list with $k > 4$ lines 9,10 and 20,21 of Algorithm 3 should also allow for the coefficients of the vectors \mathbf{v}_i , \mathbf{p} and \mathbf{v}_{k-1} to take more values than ± 1 (see for example [26, Theorem 2.2.2]).

The “higher” sieving algorithms which we considered by making the generalisation from the GaussSieve towards the MinkowskiSieve will contribute towards an asymptotic computational argument. But first we state a heuristic assumption which we will use.

Heuristic Assumption 2 (Convergence of the MinkowskiSieve) *Consider the k -MinkowskiSieve algorithm equipped with the function PrimeMinkowskiReduce (Algorithm 3). Then this algorithm will converge to a set MinkowskiSieve $_k(\mathcal{L})$.*

Remark 17. We leave the investigation for a proof of this heuristic assumption as an open problem for future research.

Theorem 9 (Shape of $\text{MinkowskiSieve}_k(\mathcal{L})$). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . We consider the k -MinkowskiSieve algorithm equipped with the function $\text{PrimeMinkowskiReduce}$. Under Heuristic Assumption 2, as k increases the set $\text{MinkowskiSieve}_k(\mathcal{L})$ converges to a set $P^+(\mathcal{L})$.*

Proof. In order to simplify the proof and avoid ambiguities we make the following convention. Both sets $P^+(\mathcal{L})$ and $\text{MinkowskiSieve}_k(\mathcal{L})$ are defined/constructed in such a way that for a vector \mathbf{v} only one of $\pm\mathbf{v}$ belongs to the set. This allows many possible choices for these sets. In order to avoid this kind of ambiguities we make the convention that a vector \mathbf{v} is included in the aforementioned sets only if its first non-zero coordinate is positive.

Initially we will prove that for every $k \geq 2$ there exists a set $P^+(\mathcal{L})$ and a set $P_k(\mathcal{L})$ such that

$$P^+(\mathcal{L}) \subseteq \text{MinkowskiSieve}_k(\mathcal{L}) \subseteq P_k(\mathcal{L}). \quad (17)$$

Let $k \geq 2$ and $\text{MinkowskiSieve}_k(\mathcal{L})$ be the converging set of an execution of the k -MinkowskiSieve. As mentioned before, we can transfer Lemma 5 from the case of GaussSieve to the k -MinkowskiSieve algorithm described in this section. This implies that for every $\text{MinkowskiSieve}_k(\mathcal{L})$ there will exist a set $P^+(\mathcal{L})$ such that $P^+(\mathcal{L}) \subseteq \text{MinkowskiSieve}_k(\mathcal{L})$. We fix this set $P^+(\mathcal{L})$.

Let $\mathbf{v} \in \text{MinkowskiSieve}_k(\mathcal{L})$ and $\mathbf{p}_1, \dots, \mathbf{p}_{k-1} \in P^+(\mathcal{L})$ with $\|\mathbf{p}_i\| < \|\mathbf{v}\|$. As the set $\text{MinkowskiSieve}_k(\mathcal{L})$ is k -reduced according to the notion implied by Algorithm 3 we can conclude that \mathbf{v} cannot be reduced by any vector of the form $\sum_{i=1}^l (-1)^{a_i} \mathbf{p}_i$ for $1 \leq l \leq k-1$. As the vectors $\mathbf{p}_1, \dots, \mathbf{p}_{k-1}$ belong to the set $\text{MinkowskiSieve}_k(\mathcal{L})$ as well, they are $k-1$ -reduced. This in turn implies that the vectors of the form $\sum_{i=1}^l (-1)^{a_i} \mathbf{p}_i$ belong to the set $P^{(l)}(\mathcal{L})$ for $1 \leq l \leq k-1$. This holds for any tuple of l vectors in $P^+(\mathcal{L})$. Hence, the set of vectors emerging from the union of all $\{\sum_{i=1}^l (-1)^{a_i} \mathbf{p}_i\}$ will be exactly $P^{(l)}(\mathcal{L})$. This implies that \mathbf{v} cannot be reduced by any vector in $\cup_{i=1}^{k-1} P^{(i)}(\mathcal{L})$. This is equivalent to the condition a vector \mathbf{v} has to satisfy according to definition 12 in order to belong to $P_k(\mathcal{L})$. Thus we can conclude that $\text{MinkowskiSieve}_k(\mathcal{L})$ is included in the $P_k(\mathcal{L})$ implied by the set $P(\mathcal{L}) = P^+(\mathcal{L}) \cup (-P^+(\mathcal{L}))$. This concludes the first part of the proof.

For the second part of the proof we distinguish between the cases of $\text{Irr}(\mathcal{L}) = P(\mathcal{L})$ and $\text{Irr}(\mathcal{L}) \neq P(\mathcal{L})$.

If it holds that $\text{Irr}(\mathcal{L}) = P(\mathcal{L})$ then apart from $P(\mathcal{L})$ being uniquely determined the same holds for the sets $P_k(\mathcal{L})$. Hence, for every $k \geq 2$ the boundary sets in (17) are uniquely determined. This enables a direct use of Lemma 6. As k increases the set $\text{MinkowskiSieve}_k(\mathcal{L})$ will be contained to even smaller and smaller sets $P_k(\mathcal{L})$ which converge to $\text{Irr}(\mathcal{L})$ according to (i) and (iii) of Lemma 6. Therefore for the limit case it could be stated that

$$P^+(\mathcal{L}) \subseteq \lim_{k \rightarrow \infty} \text{MinkowskiSieve}_k(\mathcal{L}) \subseteq \text{Irr}(\mathcal{L}). \quad (18)$$

But we assumed $\text{Irr}(\mathcal{L}) = \text{P}(\mathcal{L})$ and thus we can conclude that

$$\lim_{k \rightarrow \infty} \text{MinkowskiSieve}_k(\mathcal{L}) = \text{P}^+(\mathcal{L}).$$

In order to finish the proof we have to deal with the case $\text{Irr}(\mathcal{L}) \neq \text{P}(\mathcal{L})$. In this case, the sets $\text{P}^+(\mathcal{L})$ and $\text{P}_k(\mathcal{L})$ used in inequality (17) are not uniquely determined and therefore Lemma 6 cannot be used directly. In Lemma 6 it was shown that given the sequence of $\text{P}_k(\mathcal{L})$ implied by any $\text{P}(\mathcal{L})$ then $\lim_{k \rightarrow \infty} \text{P}_k(\mathcal{L}) = \text{Irr}(\mathcal{L})$. Hence any $\text{P}_k(\mathcal{L})$ belongs to a sequence converging to the same limit, $\text{Irr}(\mathcal{L})$. Interchanging terms ($\text{P}_k(\mathcal{L})$) among these sequences does not affect their limit. Therefore, we can again use inequality (17) and “take limits” leading to a result like (18). We have to be careful though. The right hand-side limit (i.e. $\text{Irr}(\mathcal{L})$) is well-defined but the left one can cycle over all choices of $\text{P}^+(\mathcal{L})$. This is expected as the limit of the sequence $\text{MinkowskiSieve}_k(\mathcal{L})$ as $k \rightarrow \infty$ is not unique but depends on the choice of representatives made for each non-trivial class of vectors. For convenience we assume that $\forall k > k_0$ for some k_0 this choice stabilises to some random but fixed choice. Thus, we have again reached inequality (18).

We examine the sets in inequality (18) according to the Gauss-reduced property. Let $k \geq 2$, the set $\text{MinkowskiSieve}_k(\mathcal{L})$ is a set in which the algorithm converges to and also possesses the Gauss-reduced property by construction. This holds for every $k \geq 2$ and thus transfers to the limit as well, as $k \rightarrow \infty$. The set $\text{P}^+(\mathcal{L})$ is not a maximal subset of \mathcal{L} satisfying the Gauss-reduced property but due to its construction it is maximal in the set $\text{Irr}(\mathcal{L})$. Hence, inequality (18) and maximality of $\text{P}^+(\mathcal{L})$ in $\text{Irr}(\mathcal{L})$ imply the result. \square

The conclusion in Theorem 9 is supported by the experimental results given in Table 1.

Remark 18. Theorem 9 describes asymptotic behaviour of the modified MinkowskiSieve algorithm with the goal of providing a faster way of computing sets $\text{P}^+(\mathcal{L})$. Even though, asymptotically, the algorithm possesses the desired behaviour, this does not make it immediately a computational tool for $\text{P}^+(\mathcal{L})$. There are two obstacles towards that goal. The first one is, given a lattice \mathcal{L} in dimension n , to find for which $k \geq 2$ to run k - MinkowskiSieve . This k should not be too high in order to be computationally efficient to run the algorithm. The second problem is finding for how long this k - MinkowskiSieve should run in order to approximate well enough a set $\text{MinkowskiSieve}_k(\mathcal{L})$.

5.3 Experimental results

In this section we provide some experimental results which support our claims in the previous subsections. In particular, as a first step we computed the sets $\text{R}(\mathcal{L})$, $\text{Irr}(\mathcal{L})$, $\text{P}(\mathcal{L})$ for 10 lattices in dimension 20 and afterwards we computed the output of the GaussSieve , the Triple and the $\text{Quadruple MinkowskiSieve}$. In order to generate 10 lattices in dimension 20 we used the Sage computer algebra system [33]. In particular we used Sage’s “Hard lattice generator” with the following choice of parameters,

```
sage.crypto.gen_lattice(type='random', n=1, m=20, q=10^42, seed=seed)
```

and 10 different values of *seed*. Initially, using the OpenMP parallel implementation build for the projects [5,12] we computed the set of relevant vectors $R(\mathcal{L})$ for each lattice. On top of this code (which the authors were so kind to provide us) we implemented the method described in Section 5.1 and computed the set $\text{Irr}(\mathcal{L})$. As for our experiments the lattices used were generated randomly, they did not possess any specific structure and hence $P(\mathcal{L}) = \text{Irr}(\mathcal{L})$ for all of them. This part of the experiments was performed on a node of the Lisa cluster [31] with a 16-core CPU (2.10GHz) and 96 GB of RAM. The computation of the sets $R(\mathcal{L})$ and $\text{Irr}(\mathcal{L})$ using the aforementioned implementation and hardware took about 5.5 seconds per lattice.

Finally, by modifying the already existing sieve implementations in FPLLL [32] we computed the output of the GaussSieve, Triple and Quadruple MinkowskiSieve as described in Section 5.2 for the same 10 lattices. The modifications which we made to the already existing FPLLL implementations were:

- A vector is allowed to be reduced only by a shorter vector.
- The termination condition is changed to a fixed number of collisions: $5 \cdot 10^5$ for the GaussSieve and 10^5 for the Triple and Quadruple MinkowskiSieve. These numbers were chosen to ensure the created list by the algorithm remains unchanged for “many” iterations before the algorithm terminates. These choices seem to not be optimal according to our experimental data and could possibly be further reduced.

This part of the experiments was performed on a Lenovo X250 laptop with 4 Intel Core i3-5010U CPU (2.10GHz) and 8 GB of RAM. The output of these experiments is summarised in Table 1.

Table 1 motivates a number of remarks about the involved sets. Initially, the number of relevant vectors observed was indeed close to the expected number $2 \cdot (2^{20} - 1)$. Also, the sets $\text{Irr}(\mathcal{L})$ and $P(\mathcal{L})$ were equal in all 10 cases, as we had assumed for random lattices without any underlying structure. The size of $P(\mathcal{L})$ (and $\text{Irr}(\mathcal{L})$ in this case) was observed to be some orders of magnitude smaller than the size of $R(\mathcal{L})$ making it more appealing to use in practice.

The right part of Table 1 justifies our idea to try and correlate the output of sieving algorithms with the set of irreducible vectors. Even though we cannot display here the lists of vectors which we computed but rather only their sizes, we observed the following behaviour. The list of vectors outputted by the GaussSieve contained the set $P(\mathcal{L})$ in 8 out of the 10 cases and in the other two of them there was only 1 vector missing. This supports our claim that the GaussSieve converges to a superset of $P(\mathcal{L})$. Also, as we moved to “higher” sieving algorithms like our modified version of the Triple and Quadruple MinkowskiSieve the output of the sieving algorithms approximated even closer the set $P(\mathcal{L})$. Actually, it is not a coincidence that the numbers in the columns “4-red” and “ $|\text{Irr}(\mathcal{L})|, |P(\mathcal{L})|$ ” in Table 1 differ only by a factor of 2. The output of the Quadruple MinkowskiSieve in all 10 cases gave exactly a set $P^+(\mathcal{L})$ as for every vector \mathbf{v} it stores only one of $\pm \mathbf{v}$.

Table 1: The following tables describe the sizes of the lists involved in our experiments with 10 random lattices in dimension 20. The first columns indicate the seed used for the generation of the lattice. The table on the left gives the sizes of the corresponding sets $R(\mathcal{L})$, $Irr(\mathcal{L})$ and $P(\mathcal{L})$ for each lattice. The factor 2 is due to the sign symmetry. The table on the right shows the sizes of the lists generated by the modified GaussSieve, Triple and Quadruple MinkowskiSieve.

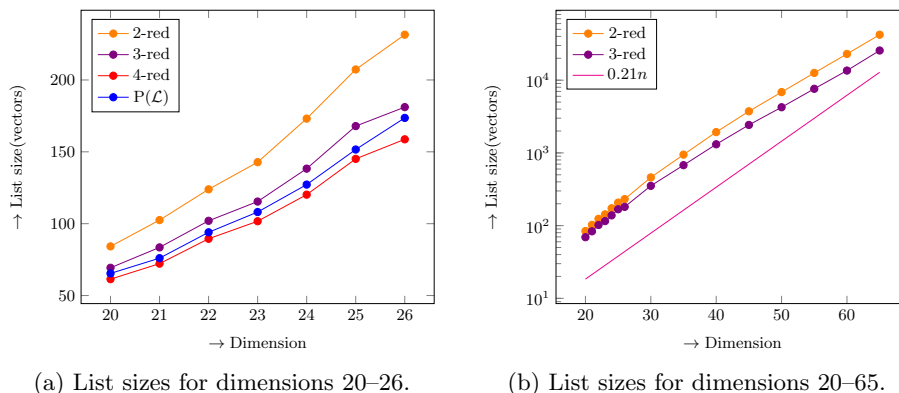
Seed	$ R(\mathcal{L}) $	$ Irr(\mathcal{L}) , P(\mathcal{L}) $	Seed	2-red	3-red	4-red
314	$2 \cdot 1048361$	$2 \cdot 66$	314	86	77	66
417	$2 \cdot 1048388$	$2 \cdot 70$	417	95	80	70
849	$2 \cdot 1048389$	$2 \cdot 68$	849	98	85	68
422	$2 \cdot 1048349$	$2 \cdot 67$	422	93	74	67
168	$2 \cdot 1048371$	$2 \cdot 60$	168	88	69	60
84	$2 \cdot 1048363$	$2 \cdot 64$	84	92	75	64
105	$2 \cdot 1048375$	$2 \cdot 62$	105	88	74	62
273	$2 \cdot 1048360$	$2 \cdot 60$	273	83	68	60
390	$2 \cdot 1048376$	$2 \cdot 66$	390	89	76	66
656	$2 \cdot 1048372$	$2 \cdot 71$	656	95	79	71

Another question which could be investigated experimentally is how the expected size of $P(\mathcal{L})$ behaves as the dimension of \mathcal{L} increases. In order to develop an intuition about this behaviour we performed a number of experiments in dimensions 20–65, the results of which are shown in Figures 3a and 3b. Likewise in our experiments in dimension 20 we used the modified OpenMP parallel implementation from [5, 12] and the modified sieve implementations in FPLLL [32]. For each dimension we depict the average value amongst 10 lattices. However, as in this case we dealt with higher dimensions we reduced the number of collisions in the termination condition of the sieve algorithms to

- GaussSieve: 10,000 collisions
- Triple MinkowskiSieve: 2,500 collisions
- Quadruple MinkowskiSieve: 2,000 collisions.

Therefore the results in Figures 3a and 3b related to sieving algorithms should only be interpreted as approximations of the algorithm’s converging set size. As we will discuss later, estimating the accuracy of this approximation is left for future research. Figure 3a illustrates the result of our experiments in dimensions 20-26. We believe that for these “smaller” dimensions the approximations are “more” accurate and that is why we show them separately. Another reason is that running the OpenMP Voronoi implementation beyond these dimensions has a substantial memory requirement (tens of GB).

Computing a least squares fit for the points in the blue curve (which indicates the correct expected values for $|P(\mathcal{L})|$ under assumption 1) gives the formula $2^{0.237n+1.286}$ which closely matches the heuristic expectation for the size of $P(\mathcal{L})$, namely $2^{0.21n}$. Furthermore Figure 3a reveals that the GaussSieve gives



(a) List sizes for dimensions 20–26.

(b) List sizes for dimensions 20–65.

Fig. 3: Experimental results on the scaling of size of $P(\mathcal{L})$ according to the dimension of \mathcal{L} . Each point in the graphs corresponds to the average value taken amongst 10 lattices. The labels k -red are used to indicate the output of the modified sieve algorithms described in this work and not the ones in the literature [3, 23].

only a superset of $P(\mathcal{L})$ even for small dimensions. The Triple and Quadruple MinkowskiSieve are much closer to the blue curve. The difference between the Triple and Quadruple MinkowskiSieve is that the one lies above the blue curve and the other below it. As we already observed in Table 1 the Triple MinkowskiSieve will probably remain above it. However the Quadruple MinkowskiSieve possess the potential to reach the “correct” curve asymptotically. Of course this could also be far from the truth for higher dimensions.

In order to put these curves more into perspective we created Figure 3b which shows the average output sizes of the GaussSieve and Triple MinkowskiSieve for dimensions 20–65. We did not draw the curve of the Quadruple MinkowskiSieve as it also turns out to be quite time costly for dimensions higher than 30. At this point we must emphasise that the used modified sieving algorithms take more time in order to terminate due to the modifications which aim not in solving SVP but computing close approximations of $P(\mathcal{L})$. For instance the modified Triple MinkowskiSieve in dimension 65 took on average 3 days in order to terminate for each lattice. However this is only the average observed time. Actually one of the ten lattices used proved to be an “easier case”, terminating in under 2 hours.

Even though these results provide some intuition on what kind of relation it could be expected between the set of irreducible vectors and sieving algorithms, they also imply some questions.

A first question which would be interesting is examining the termination condition for the sieving algorithm. In our experiments we made a specific choice on the number of collisions but this was done by trial and error and could be possibly improved. In other words, we ask for a termination condition, which if it is satisfied by a sieving algorithm (as used in this section) it guarantees that

the algorithm has reached a list of vectors which cannot be further modified by the algorithm.

A second question that arises is up to what level of sieving we should get in order to either get exactly a set $P(\mathcal{L})$ or a “very good” approximation of it. In this case the Quadruple MinkowskiSieve was enough, but this might not be the case for higher dimensional lattices. Thus it would be interesting to know how does this index increase according to the dimension. So, given some termination condition, how close can a sieving algorithm approximate a set $P(\mathcal{L})$?

If these questions receive an answer it will help in making sieving algorithms a way to either compute exactly or approximately a set $P(\mathcal{L})$ of a lattice \mathcal{L} . This would be very interesting as it will provide a way to compute a set $P(\mathcal{L})$ (exactly or approximately) without having to compute the set $R(\mathcal{L})$ which is a very costly computation.

6 Applications of $P(\mathcal{L})$

Even though the sets $\text{Irr}(\mathcal{L})$ and $P(\mathcal{L})$ might be of interest in their own, examining their relation to already existing lattice problems and algorithms is a natural question that arises. We choose to focus on the set $P(\mathcal{L})$ as it seems to be the easier to compute/approximate with existing lattice algorithms.

6.1 $P(\mathcal{L})$ in the study of shortest vector(s) problems

The results in Section 4 provide some interesting conclusions about the relation of the set $P(\mathcal{L})$ to well known lattice problems. A first observation in Section 4.2 was that $S_1(\mathcal{L})$ is included in $P(\mathcal{L})$. This leads to the following result.

Proposition 10 (Finding $P(\mathcal{L})$ implies solving SVP). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . Computing a set $P(\mathcal{L})$ provides a solution to the SVP and the kissing number problem.*

The relation $S_1(\mathcal{L}) \subseteq P(\mathcal{L})$, implies that two classic lattice problems can be solved given a $P(\mathcal{L})$. Of course this holds for any superset of $P(\mathcal{L})$ as well. We combine this observation with the inclusion $P^+(\mathcal{L}) \subseteq \text{MinkowskiSieve}_k(\mathcal{L})$ for $k \geq 2$ shown in the proof of Theorem 9. This provides some extra (heuristic) evidence that some sieving algorithms will indeed output a solution to SVP or the kissing number problem if they run long enough. This is no surprise as sieving algorithms were devised for solving SVP.

Examining the relation of SIVP to the set $P(\mathcal{L})$ is probably a more interesting question. By Corollary 1 we know that for every $i = 1, \dots, n$ there exists a vector $\mathbf{v} \in \text{Irr}(\mathcal{L})$ such that $\|\mathbf{v}\| = \lambda_i(\mathcal{L})$. The following proposition completes this result.

Proposition 11 (Finding $P(\mathcal{L})$ implies solving SIVP). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n . Computing a set $P(\mathcal{L})$ provides a solution to the SIVP.*

Proof. Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be a set of linearly independent vectors in \mathcal{L} such that $\|\mathbf{v}_i\| = \lambda_i(\mathcal{L})$ for $i = 1, \dots, n$. We distinguish two cases.

Case 1: $\nexists i \geq 2$ such that $\lambda_1(\mathcal{L}) \leq \lambda_{i-1}(\mathcal{L}) < \lambda_i(\mathcal{L}) = \lambda_{i+1}(\mathcal{L})$. This implies that there exists a $k \geq 1$ such that

$$\lambda_1(\mathcal{L}) = \dots = \lambda_k(\mathcal{L}) < \lambda_{k+1}(\mathcal{L}) < \dots < \lambda_n(\mathcal{L}).$$

Then by $S_1(\mathcal{L}) \subseteq P(\mathcal{L})$ it follows that $\mathbf{v}_1, \dots, \mathbf{v}_k$ belong to $P(\mathcal{L})$. In addition, by Corollary 1 and the definition of $P(\mathcal{L})$ it follows that all the $\mathbf{v}_{k+1}, \dots, \mathbf{v}_n$ will be included in $P(\mathcal{L})$.

Case 2: $\exists i \geq 2$ such that $\lambda_1(\mathcal{L}) \leq \lambda_{i-1}(\mathcal{L}) < \lambda_i(\mathcal{L}) = \lambda_{i+1}(\mathcal{L})$. Let $i \geq 2$ such that the condition holds. We set $k = \max\{j > i \mid \lambda_i(\mathcal{L}) = \lambda_j(\mathcal{L})\}$. Hence,

$$\lambda_1(\mathcal{L}) \leq \lambda_{i-1}(\mathcal{L}) < \lambda_i(\mathcal{L}) = \lambda_{i+1}(\mathcal{L}) = \dots = \lambda_k(\mathcal{L}).$$

We will show that $\mathbf{v}_i, \dots, \mathbf{v}_k \in P(\mathcal{L})$. Let $j \in \{i, \dots, k\}$ we set \mathcal{L}_{λ_j} to be the sublattice of \mathcal{L} spanned by all the vectors in \mathcal{L} strictly shorter than λ_j . As $\lambda_i(\mathcal{L}) = \lambda_j(\mathcal{L})$ it follows that $\mathcal{L}_{\lambda_j} = \mathcal{L}_{\lambda_i}$ which has rank $i - 1$. Assume that $\mathbf{v}_j \in \mathcal{L}_{\lambda_j}$. Then we would get that the set $\{\mathbf{v}_1, \dots, \mathbf{v}_{i-1}, \mathbf{v}_j\}$ is a set of linearly dependent vectors. Contradiction. Thus $\mathbf{v}_j \notin \mathcal{L}_{\lambda_j}$ and by Proposition 4 we get that $\mathbf{v}_j \in \text{Irr}(\mathcal{L})$. This holds for any $i \leq j \leq k$ and therefore we get that all \mathbf{v}_j with $i \leq j \leq k$ belong to $\text{Irr}(\mathcal{L})$. In order to show that they also do belong to a $P(\mathcal{L})$ it suffices to show that for every μ, ν such that $i \leq \mu < \nu \leq k$ it holds that $\|\mathbf{v}_\nu - \mathbf{v}_\mu\| \geq \lambda_i(\mathcal{L})$. Assume that there exist μ, ν such that $i \leq \mu < \nu \leq k$ and $\|\mathbf{v}_\nu - \mathbf{v}_\mu\| < \lambda_i(\mathcal{L})$. Then it follows that $\mathbf{v}_\nu - \mathbf{v}_\mu \in \mathcal{L}_{\lambda_i}$. The set of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_{i-1}, \mathbf{v}_\mu, \mathbf{v}_\nu\}$ is a linearly independent set and thus the same holds for $\{\mathbf{v}_1, \dots, \mathbf{v}_{i-1}, \mathbf{v}_\nu - \mathbf{v}_\mu\}$. This implies a set of i linearly independent vectors in the lattice \mathcal{L}_{λ_i} which is of rank $i - 1$, contradiction.

Concluding, let \mathbf{v}_l belong to the considered linearly independent set of vectors achieving the successive minima. If $\|\mathbf{v}_l\| = \|\mathbf{v}_{l+1}\|$ or $\|\mathbf{v}_l\| = \|\mathbf{v}_{l-1}\|$ then $\mathbf{v}_l \in P(\mathcal{L})$ by the proof in “case 2”. If $\|\mathbf{v}_{l-1}\| < \|\mathbf{v}_l\| < \|\mathbf{v}_{l+1}\|$ then $\mathbf{v}_l \in P(\mathcal{L})$ by the same argument used in “case 1”. \square

Remark 19. Obtaining a set of the shortest vector(s), given a set $P(\mathcal{L})$, amounts to scanning the entire set $P(\mathcal{L})$ a number of times. Thus, sorting $P(\mathcal{L})$ can be avoided.

6.2 Using $P(\mathcal{L})$ in CVPP algorithms

One main problem in lattice theory is the closest vector problem. A straightforward way of using the set $R(\mathcal{L})$ in order to solve CVPP was described in [30]. In that work, an algorithm called the iterative slicer is given which takes as input the set $R(\mathcal{L})$ and a target vector \mathbf{t} and outputs a closest lattice vector to \mathbf{t} (Algorithm 4). The main idea behind this algorithm is to iteratively reduce the target vector \mathbf{t} by the relevant vectors until the resulting vector \mathbf{t}' is contained in the Voronoi cell $\mathcal{V}(\mathcal{L})$ of the lattice. Once this condition is satisfied it is known

Algorithm 4 The iterative slicer [30]

Require: The set $R(\mathcal{L})$ and a target vector \mathbf{t} .**Ensure:** A vector $\mathbf{s} \in \mathcal{L}$ closest to \mathbf{t} .

```

1:  $\mathbf{t}' \leftarrow \mathbf{t}$ 
2: for every  $\mathbf{r} \in R(\mathcal{L})$  do
3:   if  $\|\mathbf{t}' \pm \mathbf{r}\| < \|\mathbf{t}'\|$  then
4:      $\mathbf{t}' \leftarrow \mathbf{t}' \pm \mathbf{r}$ 
5:     restart the for loop
6:   end if
7: end for
8:  $\mathbf{s} = \mathbf{t} - \mathbf{t}'$ 
9: return  $\mathbf{s}$ 

```

that $\mathbf{t} - \mathbf{t}'$ is a closest lattice point to \mathbf{t} . This algorithm is shown to terminate after a finite number of iterations.

Inspired by the iterative slicer, in [22] an algorithm is described to provably solve the CVPP in $\tilde{O}(2^{2n})$ -time by using the set $R(\mathcal{L})$ as the preprocessing data. The difference between Algorithm 4 and the algorithm in [22] is that the latter selects the relevant vectors in a specific order for reduction. This results in a $\tilde{O}(2^{2n})$ -time and $\tilde{O}(2^n)$ -space algorithm. This work was further improved in [4] by optimising the use of the preprocessing data.

However, using the set $R(\mathcal{L})$ in practice is not convenient due to its expected size of about $2^{n+1} - 2$ vectors. One way to reduce the memory requirements could be the use of a compact representation of $R(\mathcal{L})$ like the one described in [16]. In such a scenario a superset of $R(\mathcal{L})$ would be generated on the fly by a CVPP algorithm which would only use a smaller set of vectors in order to generate $R(\mathcal{L})$.

Another way would be to use a subset of $R(\mathcal{L})$ instead of the entire set. Such an approach was introduced in [18]. In that work an approximate Voronoi cell is defined as the cell implied by a list of short lattice vectors which is potentially a subset of the set $R(\mathcal{L})$. That lead to a heuristic algorithm for CVPP using the approach of Micciancio–Voulgaris but with more practical time and space complexities.

We describe a CVPP algorithm (the tuple slicer, Algorithm 5) using the set $P(\mathcal{L})$, and we discuss its advantages and disadvantages against already existing approaches. We distinguish two cases.

If $C = 1$ in Algorithm 5 then it just uses a subset of $R(\mathcal{L})$. In this case the analysis of the algorithm just follows under the “approximate Voronoi cell” approach where a specific choice has been made on the used subset. The advantage in this case is that it is guaranteed that the used list of vectors is a subset of $R(\mathcal{L})$.

If $C > 1$ Algorithm 5 behaves similar to the tuple sieving approach in [3]. A vector is reduced not only by a single vector but also by the sums of small tuples of vectors in the used list. Hence, a target vector \mathbf{t} is reduced by a superset of $P(\mathcal{L})$. If this superset includes the set $R(\mathcal{L})$ then [30, Lemma 5] guarantees the

Algorithm 5 The tuple slicer

Require: A set $P(\mathcal{L})$, a $C \in \mathbb{N}$ and a target vector \mathbf{t} .

Ensure: A vector $\mathbf{s} \in \mathcal{L}$ closest to \mathbf{t} .

```

1:  $\mathbf{t}' \leftarrow \mathbf{t}$ 
2: for  $l = 1$  to  $C$  do
3:   for all  $\{\mathbf{v}_1, \dots, \mathbf{v}_l\} \subset P(\mathcal{L})$  do
4:      $\mathbf{w} \leftarrow \sum_{i=1}^l \mathbf{v}_i$ 
5:     if  $\|\mathbf{t}' - \mathbf{w}\| < \|\mathbf{t}'\|$  then
6:        $\mathbf{t}' \leftarrow \mathbf{t}' - \mathbf{w}$ 
7:       restart the outer for loop
8:     end if
9:   end for
10: end for
11:  $\mathbf{s} = \mathbf{t} - \mathbf{t}'$ 
12: return  $\mathbf{s}$ 

```

correctness of the algorithm. This depends on the value of C . We can prove that there always exists a value of C which guarantees the inclusion of $R(\mathcal{L})$ in the generated superset.

Remark 20. In line 3 of Algorithm 5 it considers sets of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_l\}$ such that $\mathbf{v}_i \neq -\mathbf{v}_j$ but it could be that $\mathbf{v}_i = \mathbf{v}_j$.

Definition 13 (*k-wise sum of $P(\mathcal{L})$*). Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and k a positive integer. We define

$$kP(\mathcal{L}) = \left\{ \sum_{i=1}^j \mathbf{p}_i \mid \mathbf{p}_i \in P(\mathcal{L}) \text{ and } j = 1, \dots, k \right\}.$$

Proposition 12 (**Finding $R(\mathcal{L})$ via $kP(\mathcal{L})$**). Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $P(\mathcal{L})$ a complete system of irreducible vectors of it. Then there exists a positive integer $n_0 \in \mathbb{N}$ such that $R(\mathcal{L}) \subseteq n_0 P(\mathcal{L})$.

Proof. By Proposition 7 there exists a generating set $\mathbf{G} \subseteq P(\mathcal{L})$ with $|\mathbf{G}| = l \geq n$. Let $\mathbf{r} \in R(\mathcal{L})$, then there exists an $\mathbf{x} \in \mathbb{Z}^l$ such that $\mathbf{G}\mathbf{x} = \mathbf{r}$. With $\mathbf{x} = (x_1, \dots, x_l)$ set $m_{\mathbf{r}} = \|\mathbf{x}\|_1 = \sum_{i=1}^l |x_i|$. Then $\mathbf{r} \in m_{\mathbf{r}} P(\mathcal{L})$. Set $m = \max_{\mathbf{r} \in R(\mathcal{L})} \{m_{\mathbf{r}}\}$. As $R(\mathcal{L})$ is finite then m is finite and $\forall \mathbf{r} \in R(\mathcal{L})$ it holds $\mathbf{r} \in m P(\mathcal{L})$. \square

The used superset is computed on the fly. This allows for a time-memory trade-off. The algorithm loses on time complexity as it examines a larger list of vectors but it gains on the memory requirement as it stores a provably smaller subset of $R(\mathcal{L})$. In more detail the space complexity of the algorithm is proportional to $|P(\mathcal{L})|$ which can be bounded by $O(\tau_n)$. The time complexity will depend on the size of $P(\mathcal{L})$ but also on the parameter C . Following the analysis of [22] we can argue that the time complexity of Algorithm 5 will be $O(|P(\mathcal{L})|^C \cdot 2^n \text{poly}(n))$.




Fig. 4: Preliminary experimental results on the success probability of Algorithm 5. The algorithm was tested on lattices of dimensions 20, 21, 22, 23, 24. For each dimension the algorithm was tested with input $C = 1, 2, 3$ against 10000 CVP instances. Each of the 10000 CVP blocks was formed by 10 smaller blocks of 1000 CVPs corresponding to 10 lattices. Each point in the graph corresponds to the ratio of correct answers out of the 10000 CVP instances.

Remark 21. From Theorem 12 it follows that if Algorithm 5 was to be applied to the lattice family A_n^* , it should consider a value of C as high as $(n + 1)/2$ in order for $R(A_n^*)$ to be included in the used superset. Therefore, a provable upper bound on C alone will not lead to any good bound for the time complexity of Algorithm 5 in a provable setting.

Considering Algorithm 5 in a heuristic setting seems to be a more appealing choice. In such a scenario the requirements of the algorithm can be relaxed in mainly two directions. The first one is using an approximation (a superset) of $P(\mathcal{L})$ instead of the set itself. Hence, the output of the MinkowskiSieve as described in Section 5.2 could serve as such a choice. Furthermore, choosing a specific approximation of $P(\mathcal{L})$ can allow fixing the value of the parameter C in the following way.

By a heuristic result of [18] we know that if a list L containing $2^{n/2+o(n)}$ lattice vectors of norm less than $\sqrt{2}\lambda_1(\mathcal{L})$ is used as input to the iterative slicer then the success probability of the algorithm is close to 1. Following this guideline, a value for the parameter C can be chosen in a way that guarantees that the set of all vectors used for reduction in Algorithm 5 contains a list of $2^{n/2+o(n)}$ shortest lattice vectors.

Further options can be examined if it is allowed for the used slicing algorithm to succeed with probability much smaller than 1. In such a case the results

in [9, 10] provide a way of relating the success probability to size of the used preprocessed list and hence in our case C .

We briefly experimented on the relation of the success probability of Algorithm 5 and the parameter C . The results can be found in Figure 4. From these results we get a first indication that the success probability of Algorithm 5 increases as the value of C increases. However, obtaining a specific guideline on how to choose a value for C remains an open question.

7 Future work

We believe that the notion of irreducibility will motivate further research on the field of lattices. Due to its close relation to lattice sieving algorithms, it could be that the set $P(\mathcal{L})$ can provide further insight on this area. An interesting question would be if the usage of the set $P(\mathcal{L})$ (under some heuristic assumptions on its size) enables the proof of an upper bound on the time complexity of the GaussSieve [23].

The implications of $P(\mathcal{L})$ in cryptanalytic attacks could be an interesting topic to investigate. The set $P(\mathcal{L})$ is expected to be affected by an underlying structure in the lattice \mathcal{L} . It can thus be expected that structured lattices end up with a smaller set $P(\mathcal{L})$ than “average-case” lattices. Many of the modern lattice-based cryptosystems possess such underlying structures and hence they could serve as interesting cases to examine from this point of view.

The set $P(\mathcal{L})$ can be used as a tool in proving a behaviour of a lattice algorithm but could also be used itself (e.g. as preprocessing data of a CVPP algorithm). In Section 5.2 modified sieving algorithms were utilised in order to show how to compute it asymptotically. But the question of how to compute $P(\mathcal{L})$ exactly or approximately in practice remains open. Such a result would also imply the ability to compute a subset of $R(\mathcal{L})$ (of heuristically exponential size) without requiring the set $R(\mathcal{L})$.

Appendix B provides some experimental evidence showing that the size of a set $P(\mathcal{L})$ could vary a lot in some cases. An “average-case” result implying that if the underlying lattice is not “special” then the size of $P(\mathcal{L})$ cannot vary a lot would be of interest. A potential tool to reaching such a result could be lattice theta functions [11]. This is due to the fact that the coefficients in a lattice’s theta function actually represent the number of lattice vectors of a specific length. Therefore this property reveals the connection to the definition of $P(\mathcal{L})$.

Acknowledgements. We thank Filipe Cabeleira, Artur Mariano and Gabriel Falcao for providing their parallel implementation computing $R(\mathcal{L})$ as described in [5], which we used for our experimental results in Section 5.3. The authors thank Daniel Dadush for helpful discussions regarding how efficiently can the set $R(\mathcal{L})$ be generated by $P(\mathcal{L})$. We also thank Noah Stephens-Davidowitz for spotting a mistake in an earlier version of the paper. Emmanouil Doulgerakis is

supported by the NWO under grant 628.001.028 (FASOR). Thijs Laarhoven is supported by a Veni grant from NWO under project number 016.Veni.192.005.

References

1. Erik Agrell, Thomas Eriksson, Alexander Vardy, and Kenneth Zeger. Closest point search in lattices. *Transactions on Information Theory*, 48(8):2201–2214, 2002.
2. Martin Albrecht, Léo Ducas, Gottfried Herold, Elena Kirshanova, Eamonn Postlethwaite, and Marc Stevens. The general sieve kernel and new records in lattice reduction. In *Proceedings of the 38th EUROCRYPT*, pages 717–746. Springer, 2019.
3. Shi Bai, Thijs Laarhoven, and Damien Stehlé. Tuple lattice sieving. *LMS Journal of Computation and Mathematics*, 19(A):146–162, 2016.
4. Nicolas Bonifas and Daniel Dadush. Short paths on the Voronoi graph and the closest vector problem with preprocessing. In *Proceedings of the 26th SODA*, pages 295–314. ACM-SIAM, 2015.
5. Filipe Cabeleira, Artur Mariano, and Gabriel Falcao. Memory-optimized Voronoi cell-based parallel kernels for the shortest vector problem on lattices. In *Proceedings of the 27th EUSIPCO*, pages 1–5. IEEE, 2019.
6. Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In *Proceedings of the 17th ASIACRYPT*, pages 1–20. Springer, 2011.
7. John H. Conway and Neil J.A. Sloane. Low-dimensional lattices. VI. Voronoi reduction of three-dimensional lattices. In *Proceedings of the Mathematical and Physical Sciences*, volume 436, pages 55–68. The Royal Society, 1992.
8. John H. Conway and Neil J.A. Sloane. *Sphere packings, lattices and groups*. Springer, 1998.
9. Emmanouil Doulgerakis, Thijs Laarhoven, and Benne de Weger. Finding closest lattice vectors using approximate Voronoi cells. In *Proceedings of the 10th PQCRYPTO*, pages 3–22. Springer, 2019.
10. Léo Ducas, Thijs Laarhoven, and Wessel van Woerden. The randomized slicer for CVPP: sharper, faster, smaller, batchier. In *Proceedings of the 23rd PKC*, pages 3–36. Springer, 2020.
11. Noam Elkies. Theta functions and weighted theta functions of euclidean lattices, with some applications. <http://people.math.harvard.edu/~elkies/aws09.pdf>, 2009.
12. Gabriel Falcao, Filipe Cabeleira, Artur Mariano, and Luis Paulo Santos. Heterogeneous implementation of a Voronoi cell-based SVP solver. *IEEE Access*, 7:127012–127023, 2019.
13. Nicolas Gama, Phong Q. Nguyen, and Oded Regev. Lattice enumeration using extreme pruning. In *Proceedings of the 29th EUROCRYPT*, pages 257–278. Springer, 2010.
14. Richard Hamming. Error detecting and error correcting codes. *The Bell System Technical Journal*, 29(2):147–160, 1950.
15. Max Hopkins. Representation-theoretic techniques for independence bounds of Cayley graphs. Bachelor thesis, 2018.
16. Christoph Hunkenschröder, Gina Reuland, and Matthias Schymura. On compact representations of Voronoi cells of lattices. In *Proceedings of the 20th IPCO*, volume 11480 of *Lecture Notes in Computer Science*, pages 261–274. Springer, 2019.
17. Grigory Kabatiansky and Vladimir Levenshtein. Bounds for packings on a sphere and in space. *Problemy Peredachi Informatsii*, 14:3–25, 1978.

18. Thijs Laarhoven. Sieving for closest lattice vectors (with preprocessing). In *Proceedings of the 23rd SAC*, pages 523–542. Springer, 2016.
19. Arjen K. Lenstra, Hendrik W. Lenstra Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.
20. Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.
21. Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*. Kluwer Academic Publishers, Boston, Massachusetts, 2002.
22. Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *Proceedings of the 42nd STOC*, pages 351–358. ACM Press, 2010.
23. Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the shortest vector problem. In *Proceedings of the 21st SODA*, pages 1468–1480. ACM-SIAM, 2010.
24. Hermann Minkowski. In *Gesammelte Abhandlungen von Hermann Minkowski*, volume 2, pages 103–121. 1911.
25. J. W. Moon and L. Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3:23–28, 1965.
26. Phong Q. Nguyen and Damien Stehlé. Low-dimensional lattice basis reduction revisited. *ACM Transactions on Algorithms*, 5(4):46:1–46:48, 2009.
27. Phong Q. Nguyen and Thomas Vidick. Sieve algorithms for the shortest vector problem are practical. *Journal of Mathematical Cryptology*, 2(2):181–207, 2008.
28. Dayanand S. Rajan and Anil M. Shende. A characterization of root lattices. *Discrete Mathematics*, 161:309–314, 1996.
29. Mosche Rosenfeld. Independent sets in regular graphs. *Israel Journal of Mathematics*, 2:262–272, 1964.
30. Naftali Sommer, Meir Feder, and Ofir Shalvi. Finding the closest lattice point by iterative slicing. *SIAM Journal of Discrete Mathematics*, 23(2):715–731, 2009.
31. SURFsara. The Lisa cluster. <https://userinfo.surfsara.nl/systems/lisa>, 2019.
32. The FPLLL development team. fplll, a lattice reduction library. Available at <https://github.com/fplll/fplll>, 2019.
33. The Sage Developers. Sagemath, the Sage Mathematics Software System. <https://www.sagemath.org>, 2019.
34. Emanuele Viterbo and Ezio Biglieri. Computing the Voronoi cell of a lattice: The diamond-cutting algorithm. *IEEE Transactions on Information Theory*, 42:161–171, 1996.
35. Georgy Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. deuxième mémoire. recherches sur les paralléloèdres primitifs. *Journal für die reine und angewandte Mathematik*, 134:198–287, 1908.

A Corner cases among $S_1(\mathcal{L})$, $\text{Irr}(\mathcal{L})$, and $R(\mathcal{L})$

In Section 4.1 we posed two questions regarding the set $\text{Irr}(\mathcal{L})$: if and when the corner cases $S_1(\mathcal{L}) \subsetneq \text{Irr}(\mathcal{L}) = R(\mathcal{L})$ and $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L}) \subsetneq R(\mathcal{L})$ are possible. In this section we will give a partial answer to these questions by examining some already known families of special lattices, the duals of the root lattices D_n and A_n (see [8, Chapter 4]).

For $n \in \mathbb{N}$ with $n \geq 5$ we write⁴:

$$\mathcal{L}_n = 2D_n^*. \quad (19)$$

Then a basis of \mathcal{L}_n is the following (see [8, p. 120]):

$$\mathbf{B}_n = \{2e_i \mid 1 \leq i \leq n-1\} \cup \{1^n\}, \quad (20)$$

where 1^n represents the all-1 vector.

Theorem 10 (Properties of \mathcal{L}_n). *For every $n \in \mathbb{N}$ with $n \geq 5$*

$$\begin{aligned} S_1(\mathcal{L}_n) &= \{\pm 2e_i \mid 1 \leq i \leq n\} \quad \text{and} \\ \text{Irr}(\mathcal{L}_n) = R(\mathcal{L}_n) &= \{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n. \end{aligned}$$

Proof. By the definition of the lattice \mathcal{L}_n it is clear that $\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n \subset \mathcal{L}_n$. We will prove this theorem in three steps.

The first step is to show that $S_1(\mathcal{L}_n) = \{\pm 2e_i \mid 1 \leq i \leq n\}$.

The second step will be to show that $R(\mathcal{L}_n) \subseteq \{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n$.

Finally in the third step we will prove that $\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n \subseteq \text{Irr}(\mathcal{L}_n)$.

These three steps imply the result as $\text{Irr}(\mathcal{L}_n) \subseteq R(\mathcal{L}_n)$.

The ‘‘defining property’’ of the lattice \mathcal{L}_n , that if $\mathbf{v} = (v_1, \dots, v_n) \in \mathcal{L}_n$ then $v_i \equiv v_j \pmod{2}$ for all $1 \leq i, j \leq n$, will be used throughout the proof.

Step 1: Obtaining that $S_1(\mathcal{L}_n) = \{\pm 2e_i \mid 1 \leq i \leq n\}$ is trivial and is left as an exercise to the reader.

Step 2: Let $\mathbf{v} \notin R(\mathcal{L}_n)$ and $\mathbf{v} \neq \mathbf{0}$. Then by Theorem 1 we know that there exists a vector $\mathbf{x} \in \mathcal{L}_n \setminus \{\mathbf{0}, \mathbf{v}\}$ such that $\langle \mathbf{v}, \mathbf{x} \rangle \geq \|\mathbf{x}\|^2$. We will prove that for every vector $\mathbf{v} \in \mathcal{L}_n \setminus (\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n \cup \{\mathbf{0}\})$ there exists a vector $\mathbf{x} \in \mathcal{L}_n \setminus \{\mathbf{0}, \mathbf{v}\}$ such that $\langle \mathbf{v}, \mathbf{x} \rangle \geq \|\mathbf{x}\|^2$. This implies the desired property $R(\mathcal{L}_n) \subseteq \{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n$.

Let $\mathbf{v} \in \mathcal{L}_n \setminus (\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n \cup \{\mathbf{0}\})$, we distinguish two cases.

Case 1: Let \mathbf{v} be such that $\mathbf{v} = (v_1, \dots, v_n)$ with $v_i \equiv 1 \pmod{2}$ for all v_i . We already showed in step 1 of the proof that the shortest vectors with odd coordinates are the $\{\pm 1\}^n$. As \mathbf{v} does not belong to this set, $|v_i| \geq 1$ for all v_i , and there exists at least one v_j such that $|v_j| \geq 3$. Consider the vector $\mathbf{x} = (\text{sign}(v_1), \dots, \text{sign}(v_n))$. This is a valid lattice vector as $\mathbf{x} \in \{\pm 1\}^n \subset \mathcal{L}_n$

⁴ We choose to work with a scaling of D_n^* as in this way we get a lattice in \mathbb{Z}^n , which is easier to work with.

and $\mathbf{x} \neq \mathbf{v}$ as $\mathbf{v} \in \mathcal{L}_n \setminus (\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n)$. We check the inner product of \mathbf{v} and \mathbf{x} .

$$\langle \mathbf{v}, \mathbf{x} \rangle = \sum_{i=1}^n \text{sign}(v_i)v_i = \sum_{i=1}^n |v_i| \geq n + 2 > n = \|\mathbf{x}\|^2$$

This proves that $\mathbf{v} \notin \text{R}(\mathcal{L}_n)$.

Case 2: Let \mathbf{v} be such that $\mathbf{v} = (v_1, \dots, v_n)$ with $v_i \equiv 0 \pmod{2}$ for all v_i . As \mathbf{v} is a non-zero vector then it has at least one non-zero coordinate, let it be v_j . Also as v_j is even we can conclude that $|v_j| \geq 2$. We consider the vector $\mathbf{x} = 2 \text{sign}(v_j)e_j$. This is a valid lattice vector as $\mathbf{x} \in \{\pm 2e_i \mid 1 \leq i \leq n\} \subset \mathcal{L}_n$ and $\mathbf{x} \neq \mathbf{v}$ as $\mathbf{v} \in \mathcal{L}_n \setminus (\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n)$. We check the inner product of \mathbf{v} and \mathbf{x} .

$$\langle \mathbf{v}, \mathbf{x} \rangle = \sum_{i=1}^n x_i v_i = 2 \text{sign}(v_j)v_j = 2|v_j| \geq 4 = \|\mathbf{x}\|^2$$

This proves that again $\mathbf{v} \notin \text{R}(\mathcal{L}_n)$ concluding the proof of the second step.

Step 3: In this step we want to prove that $\{\pm 2e_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n \subseteq \text{Irr}(\mathcal{L}_n)$. In step 1 we already showed that $S_1(\mathcal{L}_n) = \{\pm 2e_i \mid 1 \leq i \leq n\}$ and we know that $S_1(\mathcal{L}_n) \subseteq \text{Irr}(\mathcal{L}_n)$ hence, we only have to show that $\{\pm 1\}^n \subseteq \text{Irr}(\mathcal{L}_n)$. Assume that $\mathbf{v} \in \{\pm 1\}^n$ and $\mathbf{v} \notin \text{Irr}(\mathcal{L}_n)$. Thus there are two strictly shorter vectors \mathbf{v}_1 and \mathbf{v}_2 such that $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$. In step 1 of the proof we showed that the vectors in $\{\pm 1\}^n$ are the shortest ones among those with odd coordinates. Therefore as \mathbf{v}_1 and \mathbf{v}_2 are strictly shorter than \mathbf{v} then it must be that they have even coordinates. This implies that \mathbf{v} can be written as a sum of vectors with even coordinates. This is a contradiction, as a sum of even numbers is never odd. \square

As a scaling of a lattice \mathcal{L} has the same properties as \mathcal{L} we get Theorem 5 already mentioned in Section 4.1.

Theorem 11 (Properties of D_n^*). *Let $n \in \mathbb{N}$ with $n \geq 5$. Then for the lattice D_n^* it holds that $S_1(D_n^*) \subsetneq \text{Irr}(D_n^*) = \text{R}(D_n^*)$. Furthermore $|\text{Irr}(D_n^*)| = 2^n + 2n$.*

This proves that $S_1(\mathcal{L}) \subsetneq \text{Irr}(\mathcal{L}) = \text{R}(\mathcal{L})$ is possible for every dimension $n \geq 5$. In order to complete this result from this point of view we give another three lattices, one for each of the dimensions $n = 2, 3, 4$ that possess the same property.

For $n = 2, 3, 4$ we write $\mathcal{L}_2 = \mathcal{L}(\mathbf{B}_2)$, $\mathcal{L}_3 = \mathcal{L}(\mathbf{B}_3)$, $\mathcal{L}_4 = \mathcal{L}(\mathbf{B}_4)$ with $\mathbf{B}_2, \mathbf{B}_3, \mathbf{B}_4$ being

$$\mathbf{B}_2 = \begin{pmatrix} 3 & 1 \\ 0 & 1 \end{pmatrix} \quad \mathbf{B}_3 = \begin{pmatrix} 3 & 0 & 1 \\ 0 & 3 & 1 \\ 0 & 0 & 1 \end{pmatrix} \quad \mathbf{B}_4 = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 3 & 0 & 1 \\ 0 & 0 & 3 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (21)$$

We leave it to the reader to verify our claim for these three lattices.

Our next goal is to derive a similar result for the case $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L}) \subsetneq \text{R}(\mathcal{L})$. In order to do so we will use a scaling of the lattices A_n^* .

For $n \in \mathbb{N}$ with $n \geq 3$, we write

$$\mathcal{M}_n = (n+1)A_n^*. \quad (22)$$

Then a basis of \mathcal{M}_n is formed by the columns of \mathbf{B}_n (see [8, p. 115]), where

$$\mathbf{B}_n = \begin{pmatrix} -n & 1 & 1 & \cdots & 1 \\ 1 & -n & 1 & \cdots & 1 \\ 1 & 1 & -n & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & -n \\ 1 & 1 & 1 & \cdots & 1 \end{pmatrix} \quad (23)$$

is an $(n+1) \times n$ matrix.

Remark 22. By the given basis \mathbf{B}_n for \mathcal{M}_n we can immediately observe that if $\mathbf{v} = (v_1, v_2, \dots, v_{n+1}) \in \mathcal{M}_n$ then $v_i \equiv v_j \pmod{n+1}$. Additionally $\sum_{i=1}^{n+1} v_i = 0$.

Theorem 12 (Properties of \mathcal{M}_n). *For every $n \in \mathbb{N}$ with $n \geq 3$,*

$$S_1(\mathcal{M}_n) = \text{Irr}(\mathcal{M}_n) = \{\pm(-n^1, 1^n)\} \quad \text{and} \\ \text{R}(\mathcal{M}_n) = \left\{ \pm(\alpha^\beta, (-\beta)^\alpha) \mid \beta = n+1-\alpha, 1 \leq \alpha \leq \frac{n+1}{2} \right\}.$$

Proof. We set $A = \{\pm(\alpha^\beta, (-\beta)^\alpha) \mid \beta = n+1-\alpha, 1 \leq \alpha \leq (n+1)/2\}$. We will prove this theorem in four steps. The first step is to show that $S_1(\mathcal{M}_n) = \{\pm(-n^1, 1^n)\}$. The second step will be to show that $\text{R}(\mathcal{M}_n) \subseteq A$. The third step will be to show that $\text{Irr}(\mathcal{M}_n) \subseteq \{\pm(-n^1, 1^n)\}$ and finally in the fourth step we will show that $A \subseteq \text{R}(\mathcal{M}_n)$.

Step 1: The vectors $\{\pm(-n^1, 1^n)\}$ have squared length $n^2 + n$ and hence we get $\lambda_1^2(\mathcal{L}) \leq n^2 + n$. This implies that a vector achieving $\lambda_1(\mathcal{L})$ cannot have a coordinate v_j such that $|v_j| \geq n+1$. Therefore a vector achieving $\lambda_1(\mathcal{L})$ belongs to A . The squared length of a vector in A is $\beta\alpha^2 + \alpha\beta^2 = (n+1)\alpha\beta$ which minimizes for $\alpha = 1$.

Step 2: Let $\mathbf{v} \in \mathcal{M}_n \setminus (A \cup \{\mathbf{0}\})$ and write it as $\mathbf{v} = (v_1, \dots, v_{n+1})$. Then there will exist at least one coordinate of \mathbf{v} , let it be v_j , such that $|v_j| \geq n+1$. This can be proved by a contradiction argument. Assume that there was no coordinate in \mathbf{v} such that $|v_j| \geq n+1$ then it would hold that $|v_i| \leq n$ for all $1 \leq i \leq n$ and by the fact that $v_i \equiv v_j \pmod{n+1}$ we can conclude that there would be at most two possible values for $|v_i|$. But the set A contains all such vectors of the lattice, hence that would imply $\mathbf{v} \in A$, contradiction. We set \mathbf{x} to be the vector having $\text{sign}(v_j)n$ in the j -th position and $-\text{sign}(v_j)$ in all other places. This is a valid lattice vector and $\mathbf{x} \neq \mathbf{v}$ as $\mathbf{v} \in \mathcal{M}_n \setminus A$. We check the inner product of \mathbf{v} and \mathbf{x} :

$$\langle \mathbf{v}, \mathbf{x} \rangle = \sum_{i=1}^{n+1} v_i x_i = |v_j|n - \text{sign}(v_j) \sum_{\substack{i=1 \\ i \neq j}}^{n+1} v_i = |v_j|n + |v_j| \geq (n+1)^2 > \|\mathbf{x}\|^2.$$

This proves that $\mathbf{v} \notin \mathbf{R}(\mathcal{M}_n)$ concluding the proof of the second step.

Step 3: Let $\mathbf{v} \in \mathcal{M}_n \setminus (\{\pm(-n^1, 1^n)\} \cup \{\mathbf{0}\})$ and write it as $\mathbf{v} = (v_1, \dots, v_{n+1})$. Then we will show that \mathbf{v} is reducible. By step 2 of the proof we know that $\mathbf{R}(\mathcal{M}_n) \subseteq A$ and as we know that $\text{Irr}(\mathcal{M}_n) \subseteq \mathbf{R}(\mathcal{M}_n)$ we can restrict our choice to $\mathbf{v} \in A \setminus \{\pm(-n^1, 1^n)\}$. As $\mathbf{v} \in A$ we can write $\mathbf{v} = \pm(\alpha^\beta, (-\beta)^\alpha)$ with $\beta = n+1-\alpha$ for some $1 < \alpha \leq (n+1)/2$. By Lemma 4 it suffices to find a lattice vector \mathbf{x} with $\|\mathbf{x}\| < \|\mathbf{v}\|$ and such that $2\langle \mathbf{v}, \mathbf{x} \rangle > \|\mathbf{x}\|^2$. Let $\gamma = \max\{|\alpha|, |\beta|\}$ and the j -th coordinate of \mathbf{v} be such that $|v_j| = \gamma$. Consider \mathbf{x} to be the vector with $\text{sign}(v_j)n$ in the j -th position and $-\text{sign}(v_j)$ in all other places. This is a valid lattice vector and $\|\mathbf{x}\| < \|\mathbf{v}\|$ as $\mathbf{x} \in S_1(\mathcal{M}_n)$ but $\mathbf{v} \notin S_1(\mathcal{M}_n)$. Then

$$\begin{aligned} 2\langle \mathbf{v}, \mathbf{x} \rangle &= 2 \sum_{i=1}^{n+1} v_i x_i = 2 \left(|v_j|n - \text{sign}(v_j) \sum_{\substack{i=1 \\ i \neq j}}^{n+1} v_i \right) = 2(\gamma n + \gamma) \\ &= 2(n+1)\gamma \geq (n+1)^2 > \|\mathbf{x}\|^2 \end{aligned}$$

as $\gamma \geq (n+1)/2$. This proves that $\mathbf{v} \notin \text{Irr}(\mathcal{M}_n)$ and therefore $\text{Irr}(\mathcal{M}_n) \subseteq \{\pm(-n^1, 1^n)\}$.

Step 4: By [7, Theorem 3] and the fact that the vectors $(-n^1, 1^n)$ form a strictly obtuse superbasis of \mathcal{M}_n (see [7]) it follows that $A \subseteq \mathbf{R}(\mathcal{M}_n)$ and finally $\mathbf{R}(\mathcal{M}_n) = A$. \square

This implies Theorem 6 already mentioned in Section 4.1.

Theorem 13 (Properties of A_n^*). *Let $n \in \mathbb{N}$ with $n \geq 3$. Then for the lattice A_n^* it holds that $S_1(A_n^*) = \text{Irr}(A_n^*) \subsetneq \mathbf{R}(A_n^*)$. Furthermore $|\text{Irr}(A_n^*)| = 2(n+1)$.*

This proves that $S_1(\mathcal{L}) = \text{Irr}(\mathcal{L}) \subsetneq \mathbf{R}(\mathcal{L})$ is possible for every dimension $n \geq 3$. In order to complete the result from this point of view we give another lattice in dimension $n = 2$ that possess the same property: $\mathcal{M}_2 = \mathcal{L}(\mathbf{B}_2)$, where

$$\mathbf{B}_2 = \begin{pmatrix} 4 & 1 \\ 1 & 4 \end{pmatrix}. \quad (24)$$

We leave it to the reader to verify this.

B Some graph-theoretical aspects

In Section 4.2 we introduced the notion of a complete system of irreducible vectors and we gave an example of how the set $\mathbf{P}(\mathcal{L})$ can be computed. In that example the use of graph theoretical tools was demonstrated in order to compute the set $\mathbf{P}(\mathcal{L})$ given the set $\text{Irr}(\mathcal{L})$. A natural question that arises is how costly this step can be.

In order to answer this question a few graph theory definitions are necessary. Graphs will be denoted by $\Gamma = (V, E)$, where V is the set of vertices and E is the set of edges. If $e = \{u, v\} \in E$ then we say that u and v are adjacent.

Definition 14 (Independent set). *Given a simple graph $\Gamma = (V, E)$ an independent set is a subset of vertices $U \subseteq V$, such that no two vertices in U are adjacent. An independent set is maximal if no vertex can be added without violating independence. An independent set of maximum cardinality is called maximum and its cardinality is denoted by $\alpha(\Gamma)$.*

Definition 15 (Class graph). *Let \mathcal{L} be a full rank lattice in \mathbb{R}^n and $S \in \text{Irr}(\mathcal{L})/\sim$. We define $\Gamma_{\mathcal{L}}(S)$ to be the graph where the set of vertices $V = S$ and there exists an edge between $\mathbf{v}_1, \mathbf{v}_2 \in V$ iff $\|\mathbf{v}_1 - \mathbf{v}_2\| < \|\mathbf{v}_1\|$.*

Computing $\text{P}(\mathcal{L})$ out of $\text{Irr}(\mathcal{L})$ amounts to solving a maximal independence set instance in $\Gamma_{\mathcal{L}}(S)$ for every class $S \in \text{Irr}(\mathcal{L})/\sim$. Therefore the complexity of this task highly depends on the size of the equivalence classes $S \in \text{Irr}(\mathcal{L})/\sim$ and $|\text{Irr}(\mathcal{L})/\sim|$. For average-case lattices the computational step from $\text{Irr}(\mathcal{L})$ to $\text{P}(\mathcal{L})$ should almost always be trivial, i.e. $\text{P}(\mathcal{L}) = \text{Irr}(\mathcal{L})$, as for all $S \in \text{Irr}(\mathcal{L})/\sim$ it is expected that $|S| = 2$. In these cases the set $\text{P}(\mathcal{L})$ is uniquely determined.

However, in case the underlying lattice \mathcal{L} possesses any kind of structure or symmetries it is expected that there will be equivalence classes $S \in \text{Irr}(\mathcal{L})/\sim$ with $|S| > 2$. In these cases the computational task of finding a maximal independent set in the corresponding class graph is not trivial anymore. In such cases the first step is to construct the corresponding graph $\Gamma_{\mathcal{L}}(S)$, which will take time $O(|S|^2)$. Then, naively computing a maximal independent set (which should always include both $\pm\mathbf{v}$) will take time $O(|S|m)$ where m is the number of edges in $\Gamma_{\mathcal{L}}(S)$ but, there are better performing algorithms for this task [20]. If we denote by h the maximum size of a class in $\text{Irr}(\mathcal{L})/\sim$ then the time complexity of computing $\text{P}(\mathcal{L})$ out of $\text{Irr}(\mathcal{L})$ will scale as $O(h^2|\text{Irr}(\mathcal{L})|)$.

Thus if there does not exist a class S with $|S|$ exponential to the dimension n then computing $\text{P}(\mathcal{L})$ out of $\text{Irr}(\mathcal{L})$ will take time $\tilde{O}(|\text{Irr}(\mathcal{L})|)$. In practice, stumbling upon a lattice \mathcal{L} possessing a class $S \in \text{Irr}(\mathcal{L})/\sim$ where $|S|$ is exponential to the dimension n is highly unlikely as such a lattice would be extremely structured. For the sake of mathematical curiosity (and nice graph pictures) we briefly investigate such a case of lattices, namely the \mathcal{L}_n for $n \geq 5$ defined in Appendix A. For our exposition we will need the following definition.

Definition 16 (Cayley graph). *Let G be a group and $T \subseteq G$ a generating set of G . The Cayley graph of G generated by T , denoted $\text{Cay}(G, T)$ is the directed graph $\Gamma = (V, E)$ where $V = G$ and $E = \{(g, gs) \mid g \in G, s \in T\}$.*

If $T = T^{-1}$ (T is closed under inverse) then $\text{Cay}(G, T)$ is an undirected graph.

In Appendix A we saw that $\text{Irr}(\mathcal{L}_n) = \{\pm 2\mathbf{e}_i \mid 1 \leq i \leq n\} \cup \{\pm 1\}^n$. Hence $\text{Irr}(\mathcal{L}_n)$ contains two equivalence classes of sizes $2n$ and 2^n respectively. The class $S_2 := \{\pm 1\}^n$ which we will study can be viewed as the group $G = \mathbb{Z}_2^n$. Two elements of S_2 are connected if their difference is shorter than \sqrt{n} , thus it is a sum of less than $n/4$ elements from the set $\{\pm 2\mathbf{e}_i \mid 1 \leq i \leq n\}$. In turn this implies that two elements of S_2 are connected in $\Gamma_{\mathcal{L}_n}(S_2)$ if they differ by




Fig. 5: The uncoloured Cayley graph $\text{Cay}(\mathbb{Z}_2^4, \varphi(T_1(\mathbb{Z}_2^5)))$ with generating set $\varphi(T_1(\mathbb{Z}_2^5)) = \{(0, 0, 0, 1), (0, 0, 1, 0), (0, 1, 0, 0), (1, 0, 0, 0), (1, 1, 1, 1)\}$. For convenience, instead of labelling the vertices of the graph by the elements of \mathbb{Z}_2^4 , we consider the elements of \mathbb{Z}_2^4 as binary representations and assign the corresponding integer (e.g. $(1, 0, 0, 0)$ maps to 8). This graph can be used in order to compute a representative set \tilde{S}_2 for the class $S_2 = \{\pm 1\}^5$ in the $\mathcal{L}_5 = 2D_5^*$ lattice. One such set is implied by the black vertices of the graph. The graph possesses 40 maximal independent sets of cardinality 4 and 16 maximal independent sets of cardinality 5 (maximum).

a sign in less than $n/4$ of their coordinates. Thus we can now use the following observation.

$$\Gamma_{\mathcal{L}_n}(S_2) \cong \text{Cay}(G, T_{\lceil n/4 \rceil}(G)) \tag{25}$$

Where $G = \mathbb{Z}_2^n$ and $T_{\lceil n/4 \rceil}(G) := \{\mathbf{x} \in G \mid 1 \leq |\text{supp}(\mathbf{x})| < \lceil n/4 \rceil\}$ and $\text{supp}(\mathbf{x})$ denotes the support of \mathbf{x} . In our case $T_{\lceil n/4 \rceil}(G) = T_{\lceil n/4 \rceil}^{-1}(G)$ and thus $\text{Cay}(G, T_{\lceil n/4 \rceil}(G))$ is an undirected graph.⁵ We are interested in the maximal independent sets of the graph $\Gamma_{\mathcal{L}_n}(S_2)$, but not in all of them. It is additionally required that for every vector $\mathbf{v} \in \tilde{S}_2$ also $-\mathbf{v} \in \tilde{S}_2$. This could be phrased as we work “modulo sign”. This property can be translated algebraically by working in the quotient group $H = \mathbb{Z}_2^n / \langle (1, \dots, 1) \rangle$ instead of $G = \mathbb{Z}_2^n$. Using the group

⁵ Such type of Cayley graphs are of an interest in coding theory as independent sets of $\text{Cay}(\mathbb{Z}_q^n, T_d(\mathbb{Z}_q^n))$ with $T_d(\mathbb{Z}_q^n) = \{x \in \mathbb{Z}_q^n \mid 1 \leq |\text{supp}(\mathbf{x})| < d\}$ correspond to q -ary (n, d) codes.

isomorphism

$$\begin{aligned} \varphi : \mathbb{Z}_2^n / \langle (1, \dots, 1) \rangle &\rightarrow \mathbb{Z}_2^{n-1} \\ (x_i)_{i=1}^n + \langle (1, \dots, 1) \rangle &\mapsto (x_n + x_i)_{i=1}^{n-1} \end{aligned}$$

we can transfer the problem to the graph $\Gamma_{sign} = \text{Cay}(\mathbb{Z}_2^{n-1}, \varphi(T_{\lceil n/4 \rceil}(G)))$. Each independent set in Γ_{sign} implies an independent set in $\Gamma_{\mathcal{L}_n}(S_2)$ which possesses the extra property of the “sign symmetry”. A first remark regarding the set of maximal independent sets of Γ_{sign} is that it is invariant under the group action of \mathbb{Z}_2^{n-1} . For example, if we consider the graph in Figure 5, all 16 maximal independent sets of cardinality 5 can be generated by acting with \mathbb{Z}_2^4 to the given independent set formed by the black vertices.

We briefly experimented with Γ_{sign} for the first few values $n = 5, 6, 7$ in order to get a first indication of how many maximal independent sets such a graph may have and how much their size can vary.

Table 2: Using SAGE [33] we computed all possible sizes of a maximal independent set of $\text{Cay}(\mathbb{Z}_2^{n-1}, \varphi(T_{\lceil n/4 \rceil}(\mathbb{Z}_2^n)))$ for $n = 5, 6, 7$ and the corresponding frequency of these sizes.

$n = 5$			$n = 6$				
Cardinality	4	5	Cardinality	6	8	11	16
Frequency	40	16	Frequency	320	300	32	2

$n = 7$								
Cardinality	8	14	16	17	18	19	20	22
Frequency	240	1920	625548	203840	67200	13440	2800	64

As the number of maximal independent sets seems to grow super-exponentially in the dimension n we stopped at $n = 7$. Even though experimental results are useful in order to get intuition, theoretical results are those which give the final answer to a question. In our case there are some theoretical results, originating both from graph theory and coding theory which bound the sizes we experimented with.

- Let $\Gamma = (V, E)$ be a graph with $|V| = N$. In [25] it is proven that Γ can have up to $3^{N/3}$ maximal cliques in the worst case, a bound which is tight. Complementary this also proves that a graph Γ with N vertices can possess up to $3^{N/3}$ maximal independent sets in the worst case. The results in the same work also imply that the number of different sizes of maximal independent sets is upper bounded by $N - \log_2 N$ which is shown to be tight in the worst case.

- Let Γ be an m -regular graph with N vertices. In [29] it is shown that $\alpha(\Gamma)$ can be upper bounded by $\min\{\lfloor N/2 \rfloor, N - m\}$. This bound is obtainable. In the same work, a lower bound for $\alpha(\Gamma)$ is given, depending on m and N . However this bound is not uniform but depends on number theoretic properties of N, m . In our case the appropriate lower bound for $\alpha(\Gamma)$ would be $\lceil N/(m + 1) \rceil$.
- As the graph family in question, $\Gamma_{\mathcal{L}_n}(S_2)$ (and Γ_{sign}) is specific, better upper bounds can be obtained than the general ones given in [29]. This is achieved with the use of coding theory [15]. In more detail, $\alpha(\Gamma_{\mathcal{L}_n}(S_2)) = A_2(n, \lceil n/4 \rceil)$. This equality enables the use of already known upper bounds on $A_2(n, \lceil n/4 \rceil)$ from coding theory such as the Hamming bound [14]. The lower bound implied by [29] for $\alpha(\Gamma_{\mathcal{L}_n}(S_2))$ is equivalent to the Gilbert-Varshamov bound for $A_2(n, \lceil n/4 \rceil)$.