In this course we will consider mathematical objects known as *lattices*. What is a lattice? It is a set of points in n-dimensional space with a periodic structure, such as the one illustrated in Figure 1. Three dimensional lattices occur naturally in crystals, as well as in stacks of oranges. Historically, lattices were investigated since the late 18th century by mathematicians such as Lagrange, Gauss, and later Minkowski.

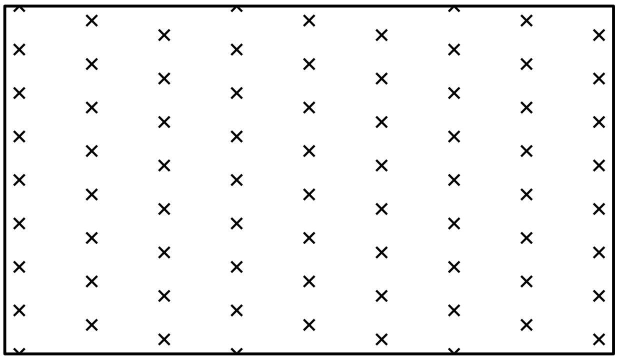

Figure 1: A lattice in $\mathbb{R}^2$

More recently, lattices have become a topic of active research in computer science. They are used as an algorithmic tool to solve a wide variety of problems; they have many applications in cryptography and cryptanalysis; and they have some unique properties from a computational complexity point of view. These are the topics that we will see in this course.

1 Lattices

We start with a more formal definition of a lattice.

DEFINITION 1 (LATTICE) *Given n linearly independent vectors $b_1, b_2, \ldots, b_n \in \mathbb{R}^m$, the lattice generated by them is defined as*

$$\mathcal{L}(b_1, b_2, \ldots, b_n) = \left\{ \sum x_i b_i \mid x_i \in \mathbb{Z} \right\}.$$

We refer to $b_1, \ldots, b_n$ as a *basis* of the lattice. Equivalently, if we define B as the $m \times n$ matrix whose columns are $b_1, b_2, \ldots, b_n$, then the lattice generated by B is

$$\mathcal{L}(B) = \mathcal{L}(b_1, b_2, \ldots, b_n) = \{Bx \mid x \in \mathbb{Z}^n\}.$$

We say that the *rank* of the lattice is n and its *dimension* is m. If $n = m$, the lattice is called a *full-rank lattice*. In this course we will usually consider full-rank lattices as the more general case is not substantially different.

Let us see some examples. The lattice generated by $(1,0)^T$ and $(0,1)^T$ is $\mathbb{Z}^2$, the lattice of all integers points (see Figure 2(a)). This basis is not unique: for example, $(1,1)^T$ and $(2,1)^T$ also generate $\mathbb{Z}^2$ (see Figure 2(b)). Yet another basis of $\mathbb{Z}^2$ is given by $(2005,1)^T, (2006,1)^T$. On the other hand, $(1,1)^T, (2,0)^T$ is not a basis of $\mathbb{Z}^2$: instead, it generates the lattice of all integer points whose coordinates sum to an even number (see Figure 2(c)). All the examples so far were of full-rank lattices. An example of a lattice that is not full is $\mathcal{L}((2,1)^T)$ (see Figure 2(d)). It is of dimension 2 and of rank 1. Finally, the lattice $\mathbb{Z} = \mathcal{L}((1))$ is a one-dimensional full-rank lattice.

DEFINITION 2 (SPAN) *The* span *of a lattice $\mathcal{L}(B)$ is the linear space spanned by its vectors,*

$$\operatorname{span}(\mathcal{L}(B)) = \operatorname{span}(B) = \{By \mid y \in \mathbb{R}^n\}.$$

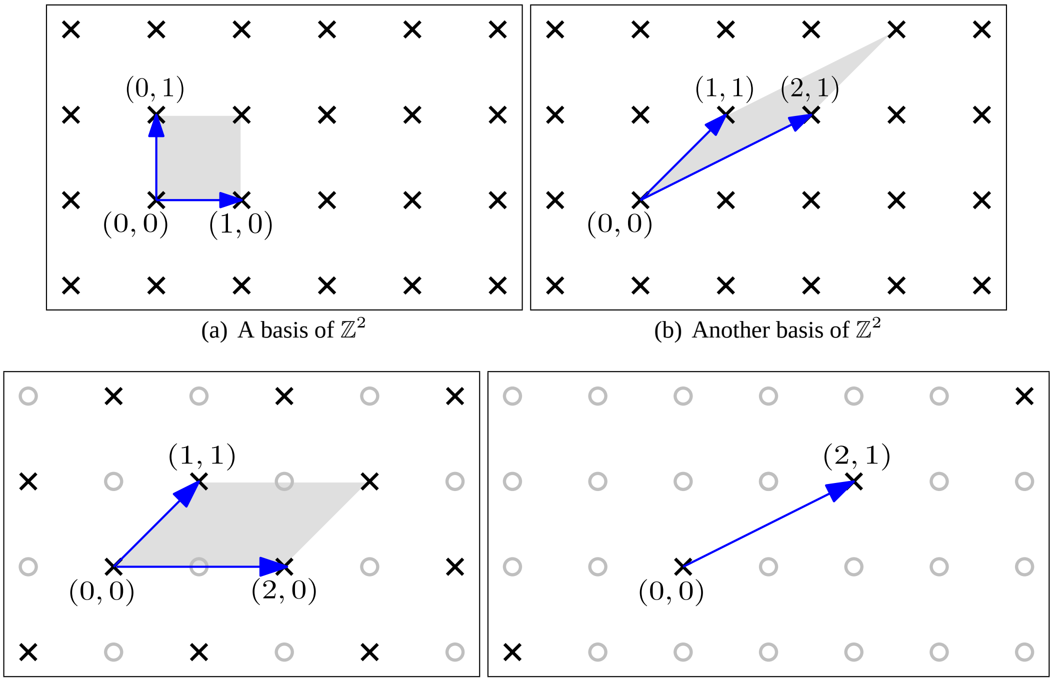

(a) A basis of $\mathbb{Z}^2$ (b) Another basis of $\mathbb{Z}^2$

(c) Not a basis of $\mathbb{Z}^2$ (d) Not a full-rank lattice

Figure 2: Some lattice bases

DEFINITION 3 (FUNDAMENTAL PARALLELEPIPED) *For any lattice basis B we define*

$$\mathcal{P}(B) = \{Bx \mid x \in \mathbb{R}^n, \forall i : 0 \leq x_i < 1\}.$$

Examples of fundamental parallelepipeds are shown by the gray areas in Figure 2. Notice that $\mathcal{P}(B)$ depends on the basis B. It follows easily from the definitions above, that if we place one copy of $\mathcal{P}(B)$ at each lattice point in $\mathcal{L}(B)$ we obtain a tiling of the entire $\mathrm{span}(\mathcal{L}(B))$. See Figure 3.

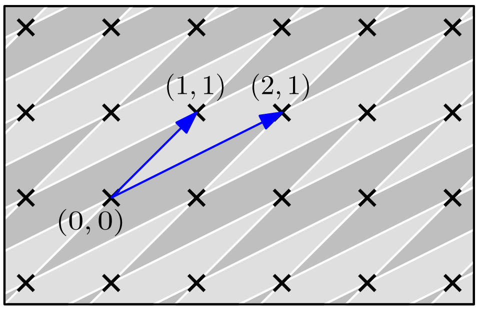

Figure 3: Tiling $\mathrm{span}(\mathcal{L}(B))$ with $\mathcal{P}(B)$

The first question we will try to answer is: how can we tell if a given set of vectors forms a basis of a lattice? As we have seen above, not every set of n linearly vectors in $\mathbb{Z}^n$ is a basis of $\mathbb{Z}^n$. One possible answer is given in the following lemma. It says that the basic parallelepiped generated by the vectors should not contain any lattice points, except the origin. As an example, notice that the basic parallelepiped shown in Figure 2(c) contains the lattice point $(1,0)$ whereas those in Figures 2(a) and 2(b) do not contain any nonzero lattice points.

LEMMA 1 *Let Λ be a lattice of rank n, and let $b_1, b_2, \ldots, b_n \in \Lambda$ be n linearly independent lattice vectors. Then $b_1, b_2, \ldots, b_n$ form a basis of Λ if and only if $\mathcal{P}(b_1, b_2, \ldots, b_n) \cap \Lambda = \{0\}$.*

PROOF: Assume first that $b_1, \ldots, b_n$ form a basis of Λ. Then, by definition, Λ is the set of all their integer combinations. Since $\mathcal{P}(b_1, \ldots, b_n)$ is defined as the set of linear combinations of $b_1, \ldots, b_n$ with coefficients in $[0, 1)$, the intersection of the two sets is $\{0\}$.

For the other direction, assume that $\mathcal{P}(b_1, b_2, \ldots, b_n) \cap \Lambda = \{0\}$. Since Λ is a rank n lattice and $b_1, \ldots, b_n$ are linearly independent, we can write any lattice vector $x \in \Lambda$ as $\sum y_i b_i$ for some $y_i \in \mathbb{R}$. Since by definition a lattice is closed under addition, the vector $x' = \sum (y_i - \lfloor y_i \rfloor) b_i$ is also in Λ. By our assumption, $x' = 0$. This implies that all y_i are integers and hence x is an integer combination of $b_1, \ldots, b_n$. $\square$

The second question we address is how to determine if two given bases B_1, B_2 are equivalent, i.e., generate the same lattice (in symbols, $\mathcal{L}(B_1) = \mathcal{L}(B_2)$). For this, we need to introduce the following definition.

DEFINITION 4 (UNIMODULAR MATRIX) *A matrix $U \in \mathbb{Z}^{n \times n}$ is called* unimodular *if* $\det U = \pm 1$.

For example, the matrix

$$\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$$

is unimodular. The following lemma appears in the homework. It tells us that the inverse of a unimodular matrix is also unimodular (so it follows that the set of unimodular matrices forms a group under matrix multiplication).

LEMMA 2 *If U unimodular, then U^{-1} is also unimodular, and in particular $U^{-1} \in \mathbb{Z}^{n \times n}$.*

LEMMA 3 *Two bases $B_1, B_2 \in \mathbb{R}^{m \times n}$ are equivalent if and only if $B_2 = B_1 U$ for some unimodular matrix U.*

PROOF: First assume that $\mathcal{L}(B_1) = \mathcal{L}(B_2)$. Then for each of the n columns b_i of B_2, $b_i \in \mathcal{L}(B_1)$. This implies that there exists an integer matrix $U \in \mathbb{Z}^{n \times n}$ for which $B_2 = B_1 U$. Similarly, there exists a $V \in \mathbb{Z}^{n \times n}$ such that $B_1 = B_2 V$. Hence $B_2 = B_1 U = B_2 V U$, and we get

$$B_2{}^T B_2 = (VU)^T B_2{}^T B_2 (VU).$$

Taking determinants, we obtain that $\det(B_2{}^T B_2) = (\det(VU))^2 \det(B_2{}^T B_2)$ and hence $\det(V)\det(U) = \pm 1$. Since V, U are both integer matrices, this means that $\det(U) = \pm 1$, as required.

For the other direction, assume that $B_2 = B_1 U$ for some unimodular matrix U. Therefore each column of B_2 is contained in $\mathcal{L}(B_1)$ and we get $\mathcal{L}(B_2) \subseteq \mathcal{L}(B_1)$. In addition, $B_1 = B_2 U^{-1}$, and since U^{-1} is unimodular (Lemma 2) we similarly get that $\mathcal{L}(B_1) \subseteq \mathcal{L}(B_2)$. We conclude that $\mathcal{L}(B_1) = \mathcal{L}(B_2)$ as required. $\square$

As an immediate corollary, we obtain that B is a basis of $\mathbb{Z}^n$ if and only if it is unimodular (verify this with the examples in Figure 2).

Another way to determine if two bases are equivalent is given in the following lemma, which is also taken from the homework.

LEMMA 4 *Two bases are equivalent if and only if one can be obtained from the other by the following operations on columns:*

1. $b_i \leftarrow b_i + kb_j$ *for some* $k \in \mathbb{Z}$*,*

2. $b_i \leftrightarrow b_j$*,*

3. $b_i \leftarrow -b_i$*.*

The last basic notion that we need is the following.

DEFINITION 5 (DETERMINANT) *Let* $\Lambda = \mathcal{L}(B)$ *be a lattice of rank* n*. We define the* determinant *of* Λ*, denoted* $\det(\Lambda)$*, as the* n*-dimensional volume of* $\mathcal{P}(B)$*. In symbols, this can be written as* $\det(\Lambda) := \sqrt{\det(B^T B)}$*. In the special case that* Λ *is a full rank lattice,* B *is a square matrix, and we have* $\det(\Lambda) = |\det(B)|$*.*

The determinant of a lattice is well-defined, in the sense that it is independent of our choice of basis B. Indeed, if B_1 and B_2 are two bases of Λ, then by Lemma 3, $B_2 = B_1 U$ for some unimodular matrix U. Hence,

$$\sqrt{\det({B_2}^T B_2)} = \sqrt{\det(U^T {B_1}^T B_1 U)} = \sqrt{\det({B_1}^T B_1)}.$$

The determinant of a lattice is inverse proportional to its density: the smaller the determinant, the denser the lattice is. In more precise terms, if one takes a large ball K (in the span of Λ) then the number of lattice points inside K approaches $\mathrm{vol}(K)/\det(\Lambda)$ as the size of K goes to infinity.

2 Gram-Schmidt Orthogonalization

Gram-Schmidt orthogonalization is a basic procedure in linear algebra that takes any set of n linearly independent vectors, and creates a set of n *orthogonal* vectors. It works by projecting each vector on the space orthogonal to the span of the previous vectors. See Figure 4 for an illustration.

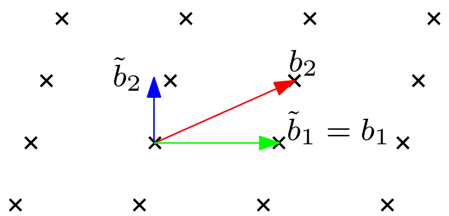

Figure 4: Gram-Schmidt orthogonalization

DEFINITION 6 *For a sequence of* n *linearly independent vectors* $b_1, b_2, \ldots, b_n$*, we define their* Gram-Schmidt orthogonalization *as the sequence of vectors* $\tilde{b}_1, \tilde{b}_2, \ldots, \tilde{b}_n$ *defined by*

$$\tilde{b}_i = b_i - \sum_{j=1}^{i-1} \mu_{i,j} \tilde{b}_j, \text{ where } \mu_{i,j} = \frac{\langle b_i, \tilde{b}_j \rangle}{\langle \tilde{b}_j, \tilde{b}_j \rangle}.$$

In words, $\tilde{b}_i$ *is the component of* b_i *orthogonal to* $\tilde{b}_1, \ldots, \tilde{b}_{i-1}$*.*

We now mention some basic and easy-to-verify properties of Gram-Schmidt orthogonalization. First, as the name suggests, for any $i \neq j$ we have that $\langle \tilde{b}_i, \tilde{b}_j \rangle = 0$. Second, for all $1 \leq i \leq n$, $\mathrm{span}(b_1, b_2, \ldots, b_i) = \mathrm{span}(\tilde{b}_1, \ldots, \tilde{b}_i)$. Third, the vectors $\tilde{b}_1, \ldots, \tilde{b}_n$ need not be a basis of $\mathcal{L}(b_1, \ldots, b_n)$. In fact, they are in general not even contained in that lattice (see Figure 4). Finally, the order of the vectors $b_1, \ldots, b_n$ matters: that is why we consider them as a sequence rather than as a set.

One useful application of the Gram-Schmidt process is the following. Let $b_1, \ldots, b_n$ be a set of n linearly independent vectors in $\mathbb{R}^m$ and consider the *orthonormal* basis given by $\tilde{b}_1/\|\tilde{b}_1\|, \ldots, \tilde{b}_n/\|\tilde{b}_n\|$. In this basis, the vectors $b_1, \ldots, b_n$ are given as the columns of the $m \times n$ matrix

$$\begin{pmatrix} \|\tilde{b}_1\| & \mu_{2,1}\|\tilde{b}_1\| & \cdots & \mu_{n,1}\|\tilde{b}_1\| \\ 0 & \|\tilde{b}_2\| & \ldots & \mu_{n,2}\|\tilde{b}_n\| \\ \vdots & & \ddots & \vdots \\ 0 & \ldots & 0 & \|\tilde{b}_n\| \\ 0 & \ldots & 0 & 0 \\ \vdots & & \vdots & \vdots \\ 0 & \ldots & 0 & 0 \end{pmatrix}. \tag{1}$$

In the case $m = n$ this is an upper-triangular square matrix. From this representation, it is easy to see that the volume of $\mathcal{P}(b_1, \ldots, b_n)$, or equivalently, $\det(\mathcal{L}(b_1, \ldots, b_n))$, is given by $\prod_{i=1}^n \|\tilde{b}_i\|$. In fact, this equality can be seen as the n-dimensional extension of the formula for computing the area of a parallelogram.

3 Successive minima

One basic parameter of a lattice is the length of the shortest nonzero vector in the lattice (we have to ask for a nonzero vector since the zero vector is always contained in a lattice and its norm is zero). This parameter is denoted by λ_1. By *length* we mean the Euclidean norm, or the ℓ_2 norm, defined as $\|x\|_2 = \sqrt{\sum x_i^2}$. We usually denote this norm simply by $\|x\|$. Occasionally in this course, we will consider other norms, such as the ℓ_1 norm, $\|x\|_1 = \sum |x_i|$ or the ℓ_∞ norm $\|x\|_\infty = \max |x_i|$.

An equivalent way to define λ_1 is the following: it is the smallest r such that the lattice points inside a ball of radius r span a space of dimension 1. This definition leads to the following generalization of λ_1, known as *successive minima*. See Figure 5.

DEFINITION 7 *Let Λ be a lattice of rank n. For $i \in \{1, \ldots, n\}$ we define the ith successive minimum as*

$$\lambda_i(\Lambda) = \inf\left\{r \mid \dim(\mathrm{span}(\Lambda \cap \overline{\mathbf{B}}(0, r))) \geq i\right\}$$

where $\overline{\mathbf{B}}(0, r) = \{x \in \mathbb{R}^m \mid \|x\| \leq r\}$ is the closed ball of radius r around 0.

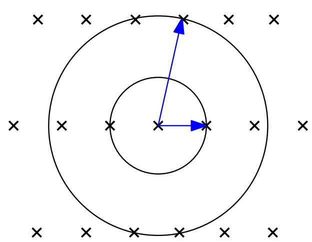

Figure 5: $\lambda_1(\Lambda) = 1, \lambda_2(\Lambda) = 2.3$

The following theorem gives a useful lower bound on the length of the shortest nonzero vector in a lattice.

THEOREM 5 *Let B be a rank-n lattice basis, and let $\tilde{B}$ be its Gram-Schmidt orthogonalization. Then*

$$\lambda_1(\mathcal{L}(B)) \geq \min_{i=1,\ldots,n} \|\tilde{b}_i\| > 0.$$

PROOF: Let $x \in \mathbb{Z}^n$ be an arbitrary nonzero integer vector, and let us show that $\|Bx\| \geq \min \|\tilde{b}_i\|$. Let $j \in \{1, \ldots, n\}$ be the largest such that $x_j \neq 0$. Then

$$|\langle Bx, \tilde{b}_j \rangle| = |\langle \textstyle\sum_{i=1}^{j} x_i b_i, \tilde{b}_j \rangle| = |x_j| \langle \tilde{b}_j, \tilde{b}_j \rangle = |x_j| \|\tilde{b}_j\|^2$$

where we used that for all $i < j$, $\langle b_i, \tilde{b}_j \rangle = 0$ and that $\langle b_j, \tilde{b}_j \rangle = \langle \tilde{b}_j, \tilde{b}_j \rangle$. On the other hand, $|\langle Bx, \tilde{b}_j \rangle| \leq \|Bx\| \cdot \|\tilde{b}_j\|$, and hence we conclude that

$$\|Bx\| \geq |x_j| \|\tilde{b}_j\| \geq \|\tilde{b}_j\| \geq \min \|\tilde{b}_i\|.$$

□

An alternative proof of Theorem 1 is the following. In the orthonormal basis $\tilde{b}_1/\|\tilde{b}_1\|, \ldots, \tilde{b}_n/\|\tilde{b}_n\|$ the lattice Λ is given by all integer combinations of the columns of the matrix in Eq. (5). It is easy to see that in any such nonzero combination, the bottom-most coordinate is at least $\min \|\tilde{b}_i\|$ in absolute value.

COROLLARY 6 *Let Λ be a lattice. Then there exists some $\varepsilon > 0$ such that $\|x - y\| > \varepsilon$ for any two non-equal lattice points $x, y \in \Lambda$.*

PROOF: For any non-equal $x, y \in \Lambda$, the vector $x - y$ is a nonzero vector in Λ. Therefore, by Theorem 5, $\|x - y\| \geq \lambda_1(\Lambda) > 0$. □

CLAIM 7 *The successive minima of a lattice are achieved, i.e., for every $1 \leq i \leq n$ there exists a vector $v_i \in \Lambda$ with $\|v_i\| = \lambda_i(\Lambda)$.*

PROOF: By Corollary 6, the ball of radius (say) $2\lambda_i(\Lambda)$ contains only finitely many lattice points. It follows from the definition of λ_i that one of these vectors must have length $\lambda_i(\Lambda)$. □

3.1 Upper bounds on the successive minima

We now present Minkowski's upper bounds on the successive minima. For simplicity, in this section we only consider full-rank lattices; it is easy to extend the results to non-full-rank lattices. We start with a theorem of Blichfeld.

THEOREM 8 (BLICHFELD) *For any full-rank lattice $\Lambda \subseteq \mathbb{R}^n$ and (measurable) set $S \subseteq \mathbb{R}^n$ with $\mathrm{vol}(S) > \det \Lambda$ there exist two nonequal points $z_1, z_2 \in S$ such that $z_1 - z_2 \in \Lambda$.*

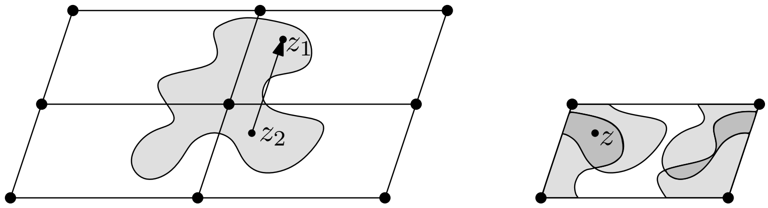

Figure 6: Blichfeldt's theorem

PROOF: Let B be a basis of Λ. As x ranges over all Λ, the sets $x+\mathcal{P}(B) := \{x+y \mid y \in \mathcal{P}(B)\}$ form a partition of $\mathbb{R}^n$. We now define $S_x = S \cap (x+\mathcal{P}(B))$ (see Figure 6). Since $S = \bigcup_{x\in\Lambda} S_x$ we conclude that $\mathrm{vol}(S) = \sum_{x\in\Lambda} \mathrm{vol}(S_x)$. We define $\widehat{S_x} = S_x - x$. Then $\widehat{S_x} \subseteq \mathcal{P}(B)$ and $\mathrm{vol}(\widehat{S_x}) = \mathrm{vol}(S_x)$, we get:

$$\sum_{x\in\Lambda} \mathrm{vol}(\widehat{S_x}) = \sum_{x\in\Lambda} \mathrm{vol}(S_x) = \mathrm{vol}(S) > \mathrm{vol}(\mathcal{P}(B)).$$

Therefore, there must exist some $x, y \in \Lambda$, $x \neq y$ for which $\widehat{S_x} \cap \widehat{S_y} \neq \emptyset$. Let z be a point in $\widehat{S_x} \cap \widehat{S_y}$. Then $z+x$ is in $S_x \subseteq S$, $z+y$ is in $S_y \subseteq S$, and $(z+x)-(z+y) = x-y$ is in Λ, as required. $\Box$

As a corollary of Blichfeld's theorem we obtain the following theorem due to Minkowski. It states that any large enough centrally-symmetric convex set contains a nonzero lattice point. A set S is *centrally-symmetric* if for any $x \in S$ we also have $-x \in S$; it is *convex* if for any $x, y \in S$ and any $\lambda \in [0,1]$ we have $\lambda x + (1-\lambda)y \in S$. It is easy to see that the theorem is false if we drop either of the two requirements on S.

THEOREM 9 (MINKOWSKI'S CONVEX BODY THEOREM) *Let Λ be a full-rank lattice of rank n. Then for any centrally-symmetric convex set S, if $\mathrm{vol}(S) > 2^n \det \Lambda$ then S contains a nonzero lattice point.*

PROOF: Define $\hat{S} = \frac{1}{2}S = \{x \mid 2x \in S\}$. Then $\mathrm{vol}(\hat{S}) = 2^{-n}\mathrm{vol}(S) > \det \Lambda$. By Theorem 8, there exist two points $z_1, z_2 \in \hat{S}$ such that $z_1 - z_2 \in \Lambda$ is a nonzero lattice point. By definition, $2z_1, 2z_2 \in S$ and because S is centrally-symmetric, also $-2z_2 \in S$. Finally, since S is convex, $\frac{2z_1 - 2z_2}{2} = z_1 - z_2$ is in S. See Figure 7. $\Box$

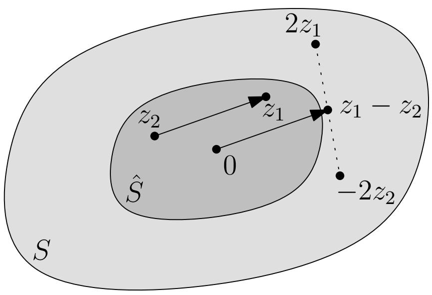

Figure 7: Minkowski's convex body theorem

CLAIM 1 *The volume of an n-dimensional ball of radius r is $\mathrm{vol}(\mathbf{B}(0,r)) \geq (\frac{2r}{\sqrt{n}})^n$.*

PROOF: This follows since this ball contains a cube of side length $\frac{2r}{\sqrt{n}}$,

$$\left\{x \in \mathbb{R}^n \mid \forall i,\ |x_i| < \frac{r}{\sqrt{n}}\right\} \subseteq \mathbf{B}(0,r).$$

$\Box$

We now obtain the following bound on the length of the shortest nonzero vector.

COROLLARY 2 (MINKOWSKI'S FIRST THEOREM) *For any full-rank lattice Λ of rank n,*

$$\lambda_1(\Lambda) \leq \sqrt{n}(\det \Lambda)^{1/n}.$$

PROOF: By definition, the (open) ball $\mathbf{B}(0, \lambda_1(\Lambda))$ contains no nonzero lattice points. By Theorem 9 and Claim 1,

$$\Big(\frac{2\lambda_1(\Lambda)}{\sqrt{n}}\Big)^n \leq \mathrm{vol}(\mathbf{B}(0, \lambda_1(\Lambda))) \leq 2^n \det \Lambda,$$

and we obtain the bound on $\lambda_1(\Lambda)$ by rearranging. $\Box$

The term $(\det \Lambda)^{1/n}$ might seem strange at first, but is in fact very natural: it makes sure that the expression scales properly. Indeed, consider the lattice $c\Lambda$ obtained by scaling Λ by a factor of c. Then clearly $\lambda_1(c\Lambda) = c\lambda_1(\Lambda)$. On the other hand, we have $\det(c\Lambda) = c^n \det(\Lambda)$, so the right hand side also scales by a factor of c, as we expect. So we could equivalently state Minkowski's first theorem as saying that any rank-n lattice with determinant 1 contains a nonzero vector of length at most $\sqrt{n}$.

How tight is this bound? It is easy to see that there are cases in which it is very far from being tight. Consider for example the lattice generated by $(\varepsilon, 0)^T$ and $(0, 1/\varepsilon)^T$ for some small $\varepsilon > 0$. Its determinant is 1 yet its shortest nonzero vector is of length ε. On the other hand, consider the lattice $\mathbb{Z}^n$. Its determinant is 1 whereas $\lambda_1(\mathbb{Z}^n) = 1$, so the bound is closer to being tight, but still not tight. In fact, it is known that for any n there exists a rank n lattice of determinant 1 whose shortest nonzero vector is of length at least $c\sqrt{n}$ for some constant c. So up to a constant, Minkowski's bound is tight. In fact, by a slightly more careful analysis, one can improve the $\sqrt{n}$ bound to $c\sqrt{n}$ for some constant $c < 1$.

Finally, we mention that in the discussion above we considered the ℓ_2 norm. It is easy to extend Minkowski's theorem to other norms. All that is required is to compute the volume of a ball under the given norm.

Minkowski's first theorem considers the shortest nonzero vector, i.e., the first successive minimum λ_1. A strengthening of the bound is given by what is known as Minkowski's second theorem. Instead of considering just λ_1, this bound considers the geometric mean of all λ_i (which is clearly at least λ_1).

THEOREM 3 (MINKOWSKI'S SECOND THEOREM) *For any full-rank lattice Λ of rank n,*

$$\Big(\prod_{i=1}^{n} \lambda_i(\Lambda)\Big)^{1/n} \leq \sqrt{n}(\det \Lambda)^{1/n}.$$

PROOF: Let $x_1, \ldots, x_n \in \Lambda$ be linearly independent vectors achieving the successive minima, $\|x_i\| = \lambda_i(\Lambda)$. Let $\tilde{x}_1, \ldots, \tilde{x}_n$ be their Gram-Schmidt orthogonalization. Consider the open ellipsoid with axes $\tilde{x}_1, \ldots, \tilde{x}_n$ and lengths $\lambda_1, \ldots, \lambda_n$,

$$T = \Big\{ y \in \mathbb{R}^n \;\Big|\; \sum_{i=1}^{n} \Big(\frac{\langle y, \tilde{x}_i \rangle}{\|\tilde{x}_i\| \cdot \lambda_i} \Big)^2 < 1 \Big\}.$$

See Figure 8.

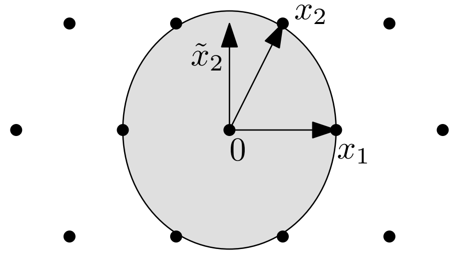

Figure 8: The ellipsoid T. The vector x_1 is on the boundary of T, and x_2 is strictly outside.

We claim that T does not contain any non-zero lattice points. Indeed, take any nonzero $y \in \Lambda$ and let $1 \leq k \leq n$ be the maximal such that $\|y\| \geq \lambda_k(\Lambda)$. It must be that $y \in \mathrm{span}(\tilde{x}_1, \dots, \tilde{x}_k) = \mathrm{span}(x_1, \dots, x_k)$, since otherwise $x_1, \dots, x_k, y$ are $k+1$ linearly independent vectors of length less than $\lambda_{k+1}(\Lambda)$. Now,

$$\sum_{i=1}^{n} \Big(\frac{\langle y, \tilde{x}_i \rangle}{\|\tilde{x}_i\| \cdot \lambda_i} \Big)^2 = \sum_{i=1}^{k} \Big(\frac{\langle y, \tilde{x}_i \rangle}{\|\tilde{x}_i\| \cdot \lambda_i} \Big)^2 \geq \frac{1}{\lambda_k^2} \sum_{i=1}^{k} \Big(\frac{\langle y, \tilde{x}_i \rangle}{\|\tilde{x}_i\|} \Big)^2 = \frac{\|y\|^2}{\lambda_k^2} \geq 1$$

and therefore, $y \notin T$.

By Minkowski's convex body theorem, $\mathrm{vol}(T) \leq 2^n \det \Lambda$. But on the other hand,

$$\mathrm{vol}(T) = \Big(\prod_{i=1}^{n} \lambda_i \Big) \mathrm{vol}(\mathbf{B}(0,1)) \geq \Big(\prod_{i=1}^{n} \lambda_i \Big) \Big(\frac{2}{\sqrt{n}} \Big)^n.$$

Combining the two bounds, we obtain that

$$\Big(\prod_{i=1}^{n} \lambda_i \Big)^{1/n} \leq \sqrt{n} (\det \Lambda)^{1/n}.$$

□

4 Computational problems

Minkowski's first theorem implies that any lattice Λ of rank n contains a nonzero vector of length at most $\sqrt{n}(\det \Lambda)^{1/n}$. Its proof, however, is non-constructive: it does not give us an algorithm to find such a lattice vector. In fact, there is no known efficient algorithm that finds such short vectors.

To discuss such computational issues, let us define the most basic computational problem involving lattices: the shortest vector problem, or SVP for short. Here, we are given a lattice and we are supposed to find the shortest nonzero lattice point. More precisely, there are three variants of the SVP, depending on whether we have to actually find the shortest vector, find its length, or just decide if it is shorter than some given number:

- **Search** SVP: Given a lattice basis $B \in \mathbb{Z}^{m \times n}$ find $v \in \mathcal{L}(B)$ such that $\|v\| = \lambda_1(\mathcal{L}(B))$.
- **Optimization** SVP: Given a lattice basis $B \in \mathbb{Z}^{m \times n}$ find $\lambda_1(\mathcal{L}(B))$.
- **Decisional** SVP: Given a lattice basis $B \in \mathbb{Z}^{m \times n}$ and a rational $r \in \mathbb{Q}$, determine whether $\lambda_1(\mathcal{L}(B) \leq r$ or not.

Notice that we restrict the lattice basis to consist of integer vectors, as opposed to arbitrary real vectors. The purpose of this is to make the input representable in finitely many bits so we can consider SVP as a standard computational problem. We could also allow the lattice basis to consist of rational vectors. This would lead to an essentially equivalent definition, since by scaling, one can make all rational coordinates integer.

Two easy relations among the three variants above is that the decision variant is not harder than the optimization variant, and that the optimization variant is not harder than the search variant. In fact, it can be shown that the converse is also true: the optimization variant is not harder than the decision variant, and the search variant is not harder than the optimization variant. To summarize, the three variants are essentially equivalent.

In this course, we will be more interested in the *approximation* variants of SVP. Here, instead of finding the shortest vector, we are interested in finding an approximation of it. The factor of approximation is given by some parameter $\gamma \geq 1$:

- **Search** SVP_γ: Given a lattice basis $B \in \mathbb{Z}^{m\times n}$ find $v \in \mathcal{L}(B)$ such that $v \neq 0$ and $\|v\| \leq \gamma \cdot \lambda_1(\mathcal{L}(B))$.
- **Optimization** SVP_γ: Given a lattice basis $B \in \mathbb{Z}^{m\times n}$ find d such that $d \leq \lambda_1(\mathcal{L}(B)) \leq \gamma \cdot d$.
- **Promise** SVP_γ: An instance of the problem is given by a pair (B, r) where $B \in \mathbb{Z}^{m\times n}$ is a lattice basis and $r \in \mathbb{Q}$. In YES instances, $\lambda_1(\mathcal{L}(B)) \leq r$. In NO instances, $\lambda_1(\mathcal{L}(B)) > \gamma \cdot r$.

The latter variant is usually denoted GapSVP_γ. It is a *promise problem*. By this we mean a problem defined by two disjoint sets of inputs: the YES instances, and the NO instances. The goal is to determine which set the input is taken from. Unlike decision problems, the union of these sets does not have to contain all possible inputs. In other words, there are illegal inputs on which the algorithm's behavior is undefined.

As before, we have that for any $\gamma \geq 1$, the promise variant is not harder than the optimization variant, and that the optimization variant is not harder than the search variant. It is also known that the optimization variant is not harder than the promise variant. Interestingly, it is an open question whether the search variant is not harder than the optimization variant.

Another fundamental lattice problem is the closest vector problem, or CVP for short. Here, the goal is to find the lattice point that is closest to a given point in space. As before, for any approximation factor $\gamma \geq 1$ we can define three variants:

- **Search** CVP_γ: Given a lattice basis $B \in \mathbb{Z}^{m\times n}$ and a vector $t \in \mathbb{Z}^m$, find $v \in \mathcal{L}(B)$ such that $\|v - t\| \leq \gamma \cdot \mathrm{dist}(t, \mathcal{L}(B))$.
- **Optimization** CVP_γ: Given a lattice basis $B \in \mathbb{Z}^{m\times n}$ and a vector $t \in \mathbb{Z}^m$, find d such that $d \leq \mathrm{dist}(t, \mathcal{L}(B)) \leq \gamma \cdot d$.
- **Promise** CVP_γ: An instance of the problem is given by a triple (B, t, r) where $B \in \mathbb{Z}^{m\times n}$ is a lattice basis, $t \in \mathbb{Z}^m$, and $r \in \mathbb{Q}$. In YES instances, $\mathrm{dist}(t, \mathcal{L}(B)) \leq r$. In NO instances, $\mathrm{dist}(t, \mathcal{L}(B)) > \gamma \cdot r$.

Both the CVP and the SVP are difficult computational problems, which we will discuss in more detail later in this course. There are also some easy computational problems involving lattices, such as:

- **Membership**: Given a lattice basis $B \in \mathbb{Z}^{m\times n}$ and a vector $v \in \mathbb{Z}^m$, decide if $v \in \mathcal{L}(B)$.
  The equation $Bx = v$ can be seen as a system of m linear equations in n variables. We can therefore solve it efficiently by Gaussian elimination. If a solution exists and it happens to be in $\mathbb{Z}^n$ (as opposed to $\mathbb{Q}^n$), output YES; otherwise output NO.
- **Equivalence**: Given $B_1, B_2 \in \mathbb{Z}^{m\times n}$, decide if $\mathcal{L}(B_1) = \mathcal{L}(B_2)$.
  To solve this, we check two things: that each column of B_1 is contained in $\mathcal{L}(B_2)$ and that each column of B_2 is contained in $\mathcal{L}(B_1)$. If $\mathcal{L}(B_1) = \mathcal{L}(B_2)$, these two checks are satisfied. Conversely, if these checks are satisfied, then $\mathcal{L}(B_1) \subseteq \mathcal{L}(B_2)$ and $\mathcal{L}(B_2) \subseteq \mathcal{L}(B_1)$ and hence $\mathcal{L}(B_1) = \mathcal{L}(B_2)$.

Several other easy computational problems are given in the homework. Vaguely speaking, what they all have in common is that they do not involve the geometry of the lattice.