Mastermath, Spring 2018
Intro to Lattice Algs & Crypto

Lecture 1

Lecturers: D. Dadush, L. Ducas
Scribe: D. Dadush

Introduction

1 Introduction

In this course, the first mathematical objects we will consider are 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.

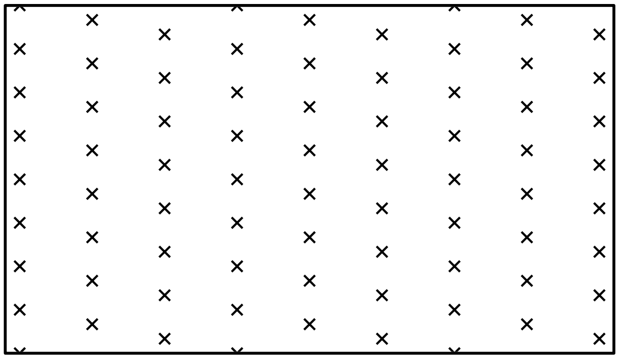

Figure 1: A lattice in $\mathbb{R}^2$

More recently, lattices have become a topic of active research in computer science. Algorithmic problems based on lattices (e.g. Shortest and Closest Vector Problems, ...) has found a wide variety of applications; they have used within optimization algorithms, in the design of wireless communication protocols, and perhaps the most active research area, in the development of secure cryptographic primitives (cryptography) and in establishing the insecurity of certain cryptographic schemes (cryptanalysis). In this course, our focus will be three-fold: to provide a solid understand of the geometric properties of lattices, to present algorithms and complexity results for fundamental lattice problems, and to explain applications of lattice techniques to the study of cryptography and cryptanalysis.

2 Notation and Basic Concepts

We use $\mathbb{R}$ for the real numbers, $\mathbb{Z}$ for the integers and $\mathbb{N}$ for the natural numbers (positive integers). Correspondly, we use $\mathbb{R}^n$ and $\mathbb{Z}^n$ to denote the n-dimensional versions for some $n \in \mathbb{N}$. We write generally matrices as $\mathbf{B}$ in uppercase bold, vectors $\mathbf{x} \in \mathbb{R}^n$ in lowercase bold and scalars as $x \in \mathbb{R}$. The Euclidean norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is denoted $\|\mathbf{x}\| \overset{\text{def}}{=} (\sum_{i=1}^n x_i^2)^{1/2}$ and the unit Euclidean ball in $\mathbb{R}^n$ is denoted by $\mathcal{B}_2^n \overset{\text{def}}{=} \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| \leq 1\}$. For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we denote their inner product by $\langle \mathbf{x}, \mathbf{y} \rangle \overset{\text{def}}{=} \sum_{i=1}^n x_i y_i$. Given a linear subspace $W \subseteq \mathbb{R}^n$, we let $W^\perp = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \forall \mathbf{y} \in W\}$ denote the orthogonal complement of W. We define $\pi_W : \mathbb{R}^n \to W$ to be the orthogonal projection onto W.

For a set $A \subseteq \mathbb{R}^n$, we use $\text{span}(A)$ to denote its linear span of A, i.e. the smallest linear subspace containing A. We define the dimension $\dim(A)$ of A to be the dimension of the linear span, that is, $\dim(A) \overset{\text{def}}{=} \dim(\text{span}(A))$. We denote the volume of A by $\text{vol}_n(A)$. For a matrix $\mathbf{T} \in \mathbb{R}^{m \times n}$, we define $\mathbf{T}A \overset{\text{def}}{=} \{\mathbf{T}\,\mathbf{a} : \mathbf{a} \in A\}$. For two sets $A, B \subseteq \mathbb{R}^n$, $s, t \in \mathbb{R}$, we define their Minkowski sum $sA + tB \overset{\text{def}}{=} \{s\mathbf{a} + t\mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B\}$.

3 Definitions of Lattices

The main goal of this lecture is to introduce the basic concept of a lattice, define one of its basic geometric parameters (the shortest non-zero vector), and present various equivalent definitions of a lattice. Our abstract definition of a lattice is given below:

Definition 1 (Lattice) *$\mathcal{L} \subseteq \mathbb{R}^n$ is a lattice if $\mathcal{L}$ is a discrete additive subgroup of $\mathbb{R}^n$. $\mathcal{L}$ is a rank k lattice, or k-dimensional, if $\dim(\mathcal{L}) = k$. $\mathcal{L}$ is said to be full-rank if $\dim(\mathcal{L}) = n$ (dimension of the ambient space).*

To be able to interpret this, one needs to define both what discrete and what additive subgroup means. Since our lattices are embedded in $\mathbb{R}^n$, the definitions given below rely on the additive structure and topology of $\mathbb{R}^n$.

1. **Discrete:** $\mathcal{L}$ is discrete if the induced topology on $\mathcal{L}$ is discrete.
   That is, every subset of $\mathcal{L}$ is open.

2. **Additive subgroup:**

   (a) $\forall \mathbf{x}, \mathbf{y} \in \mathcal{L}, \mathbf{x} + \mathbf{y} \in \mathcal{L}$.

   (b) $\forall \mathbf{x} \in \mathcal{L}, -\mathbf{x} \in \mathcal{L}$.

The above definitions encapsulate the essential properties of a lattice, however it does not tell us how to build them or describe them. As we will show, every lattice can be described either by a set of linearly independent generators, known as a basis, or via a system of modular equations. We give some examples below.

To build a rank k lattice in $\mathbb{R}^n$, we may take any set $\mathbf{b}_1, \ldots, \mathbf{b}_k \in \mathbb{R}^n$ of *linearly independent* vectors and define

$$\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_k) = \{\sum_{i=1}^{k} z_i \mathbf{b}_i : z_1, \ldots, z_k \in \mathbb{Z}\}. \qquad \text{(Basis Rep.)}$$

Here we say that $\mathbf{b}_1, \ldots, \mathbf{b}_k$ is a basis for the lattice $\mathcal{L} = \mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_k)$. We denote the basis matrix for $\mathbf{b}_1, \ldots, \mathbf{b}_k$ as $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_k)$, i.e. the matrix whose columns are $\mathbf{b}_1, \ldots, \mathbf{b}_k$. For convenience, we will use the notation $\mathcal{L}(\mathbf{B}) = \mathbf{B}\mathbb{Z}^k$, where we note that $\mathcal{L}(\mathbf{B}) = \mathcal{L}$. Furthermore, we will often interchangeably refer to $\mathbf{B}$ and $\mathbf{b}_1, \ldots, \mathbf{b}_k$ as a basis for $\mathcal{L}$. While the above definition requires linear independence, we will make an exception for the trivial lattice $\{\mathbf{0}\}$, for which we will consider $\mathbf{0}$ to be a basis.

For the most basic example, we can take $\mathbb{Z}^2 = \{(x, y) : x, y \in \mathbb{Z}\}$, i.e. the standard integer lattice in 2 dimensions. Here it is easy to see that $\mathbb{Z}^2 = \mathcal{L}(\mathbf{b}_1, \mathbf{b}_2)$ where $\mathbf{b}_1 = (0, 1)$ and $\mathbf{b}_2 = (1, 0)$. Note that $\mathbb{Z}^2$ admits more than one basis, in particular the basis $(0, 1)$, $(1, 1)$ still generates the same lattice. In fact, for any lattice $\mathcal{L}$ of rank $n > 1$ admits *infinitely* many distinct bases.

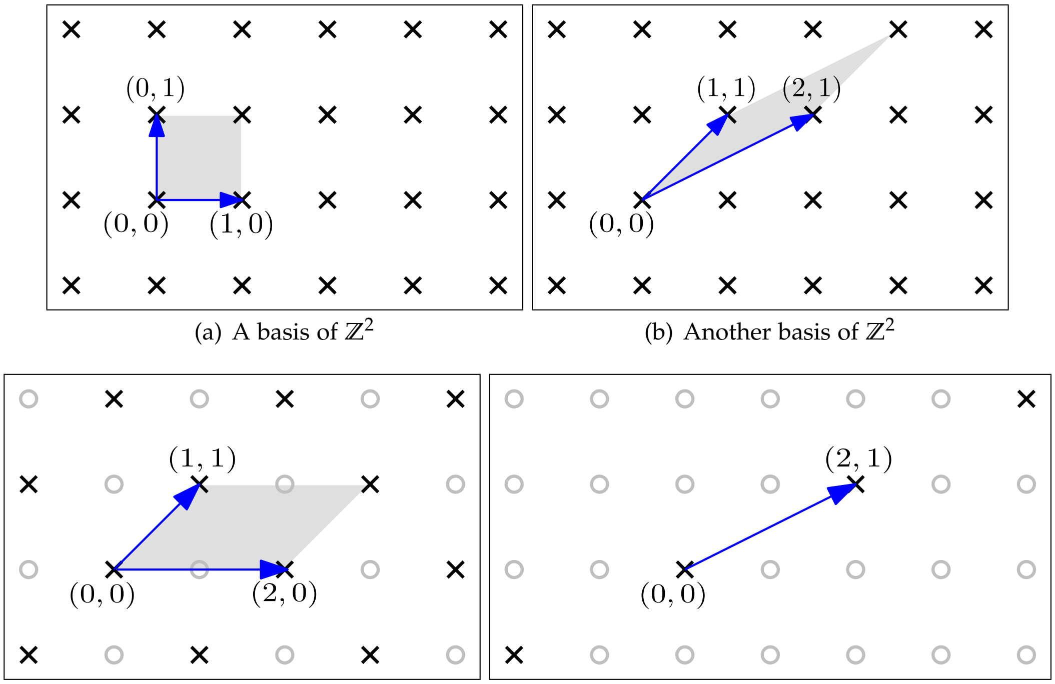

(a) A basis of $\mathbb{Z}^2$ (b) Another basis of $\mathbb{Z}^2$

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

Figure 2: Some lattice bases

We note that it is important that to use a set of linearly independent generators to build the lattice. Indeed, without this restriction one can easily build additive subgroups of $\mathbb{R}^n$ which are not lattices (e.g. the subgroup generated by $1, \sqrt{2}$ is dense in $\mathbb{R}$).

The second way to define lattices is via a set of modular equations. Given any matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we may examine:

$$\Lambda^{\perp}(\mathbf{A}) := \{\mathbf{x} \in \text{rowspan}(\mathbf{A}) : \mathbf{A}\mathbf{x} \equiv \mathbf{0} \pmod 1\} . \qquad \text{(Dual Rep.)}$$

Here, the modular equations should hold for each row, where $x \equiv y \pmod 1$ iff $x - y \in \mathbb{Z}$. Thus, the above is in fact equivalent to $\mathbf{A}\mathbf{x} \in \mathbb{Z}$. The condition $\mathbf{x} \in \text{rowspan}(A)$ is imposed to disallow non-zero vectors in the kernel of $\mathbf{A}$, since then $\Lambda^{\perp}(\mathbf{A})$ would contain an entire line through the origin and hence would not be discrete. The trivial lattice $\{0\}$ may be expressed as $\Lambda^{\perp}(\mathbf{0}^{\mathsf{T}})$, where $\mathbf{0}^{\mathsf{T}}$ corresponds to the row of zeros. Note that as long as $\Lambda^{\perp}(\mathbf{A}) \neq \{\mathbf{0}\}$, we may remove any zero row from $\mathbf{A}$ without changing the lattice. Perhaps the most important instantion of the above is to isolate a sublattice of $\mathbb{Z}^n$ via a so-called "parity check" matrix. For example, if we let C be an $m \times n$ matrix with entries in $\mathbb{Z}_p := \mathbb{Z}/p\mathbb{Z}$, then we may examine:

$$\Lambda_p^{\perp}(\mathbf{C}) = \{\mathbf{x} \in \mathbb{Z}^n : \mathbf{C}\mathbf{x} \equiv \mathbf{0} \pmod p\} . \qquad \text{(Parity Check)}$$

Lattices of the above type will be very useful in this course, as they are easy to generate at "random". In particular, one may choose the entries of $\mathbf{C}$ uniformly from $\mathbb{Z}_p$. We will see properties of these types of lattices "random lattices" later in the course. As a specific example of the above, one may look at the lattice $\{(x, y) \in \mathbb{Z}^2 : x + y \equiv 0 \pmod 2\}$, i.e. all integer points in $\mathbb{Z}^2$ whose coordinate sum is even. A little thought, will reveal that this lattice can in fact be generated by the basis $(1,1)$ and $(0,2)$, and hence this lattice can be represented using form (Basis Rep.).

While it is not obvious that every lattice of the above type admits a basis, it is easy to see that lattice of form (Parity Check) can be expressed as a lattice of form (Dual Rep.). In particular, by identifying the entries of $\mathbf{C}$ with integers in $\{0,\dots,p-1\}$, one can verify that

$$\Lambda_p^{\perp}(\mathbf{C}) = \Lambda^{\perp}\begin{pmatrix}\mathbf{C}/p \\ \mathbf{I}_n\end{pmatrix} .$$

The main theorem of this lecture is that all the different lattice representations are equivalent.

Theorem 2 (Lattice Representations) *Let $\mathcal{L} \subseteq \mathbb{R}^n$. The following are equivalent:*

1. *$\mathcal{L}$ is a lattice.*
2. *$\mathcal{L}$ can be expressed using* (Basis Rep.).
3. *$\mathcal{L}$ can be expressed using* (Dual Rep.).

We will develop the tools to prove the above theorem throughout the next sections. To begin, we first show that anything of type (Basis Rep.) or (Dual Rep.) is indeed a lattice. While the additive subgroup property is essentially immediate, discreteness (or at least useful quantitative bounds on it) requires some work and is the subject of section 4. The bulk of the effort will be in the reverse direction, where the principal aspect is showing that every lattice indeed admits a basis. This is the content of section 5.

4 The Shortest Non-Zero Vector

We begin by the main lattice parameter which controls "how discrete" a lattice is. For a lattice $\mathcal{L} \subseteq \mathbb{R}^n$, we define

$$\lambda_1(\mathcal{L}) = \inf_{\mathbf{y}\in\mathcal{L}\setminus\{\mathbf{0}\}} \|\mathbf{y}\|.$$

The following lemma establishes that the discreteness property is encapsulated by the positivity of λ_1 when the lattice is non-trivial (i.e. $\mathcal{L} \neq \{\mathbf{0}\}$).

Lemma 3 *Let $\mathcal{L}$ be a non-trivial additive subgroup of $\mathbb{R}^n$. Then the following are equivalent:*

1. *$\mathcal{L}$ is a lattice.*
2. *$\lambda_1(\mathcal{L}) > 0$.*
3. *$|\mathcal{L} \cap S| < \infty$ for any bounded set $S \subseteq \mathbb{R}^n$.*
4. *$\mathcal{L}$ contains a shortest non-zero vector.*

Proof:

$(1 \Rightarrow 2)$. Since $\mathcal{L}$ is an additive subgroup of $\mathbb{R}^n$, $\mathbf{0} \in \mathcal{L}$. By discreteness of $\mathcal{L}$, $\{\mathbf{0}\}$ forms an open set in $\mathcal{L}$. Thus, there exists a radius $r > 0$ such that $r\mathcal{B}_2^n \cap \mathcal{L} = \{\mathbf{0}\}$. Since $\mathcal{L}$ contains non-zero vectors by our non-triviality assumption, we see that $\lambda_1(\mathcal{L}) \geq r$, as needed.

(2 ⇒ 3). Let $\lambda = \lambda_1(\mathcal{L}) > 0$, and let $A = \mathcal{L} \cap r\mathcal{B}_2^n$. For any distinct $\mathbf{x}, \mathbf{y} \in A$ note that $\mathbf{x} - \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ and hence $\|\mathbf{x} - \mathbf{y}\| \geq \lambda$. From here, we must have that open balls of radius $\lambda/2$ around $\mathbf{x}$ and $\mathbf{y}$ must be interior disjoint, i.e. $(\mathbf{x} + \frac{\lambda}{2}\text{interior}(\mathcal{B}_2^n)) \cap (\mathbf{y} + \frac{\lambda}{2}\text{interior}(\mathcal{B}_2^n)) = \emptyset$. Furthermore, $\mathbf{x} + \frac{\lambda}{2}\mathcal{B}_2^n \subseteq S + \frac{\lambda}{2}\mathcal{B}_2^n$. If $|\mathcal{L} \cap S| = \infty$, then

$$\text{vol}_n(S + \frac{\lambda}{2}\mathcal{B}_2^n) \geq \sum_{\mathbf{x} \in \mathcal{L} \cap S} \text{vol}_n(\frac{\lambda}{2}\mathcal{B}_2^n) = \infty.$$

However, since $S + \frac{\lambda}{2}\mathcal{B}_2^n$ is clearly bounded (given that S is bounded), we must have that $\text{vol}_n(\frac{\lambda}{2}\mathcal{B}_2^n + S) < \infty$, a clear contradiction.

(3 ⇒ 4). Since $\mathcal{L}$ is non-trivial, we can pick $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$. Let $r = \|\mathbf{y}\|$, and examine $A = (\mathcal{L} \cap r\mathcal{B}_2^n) \setminus \{\mathbf{0}\}$. Notice that non-zero vector shorter than $\mathbf{y}$ must be contained in A. Since $r\mathcal{B}_2^n$ is bounded, by assumption we have that $|A| < \infty$. Since A is finite, we can pick $\mathbf{y} \in A$ of minimum ℓ_2 norm. By construction, $\mathbf{y}$ is a shortest-nonzero vector of $\mathcal{L}$, as needed.

(4 ⇒ 1). Let $\lambda_1 := \lambda_1(\mathcal{L})$. Since by assumption there exists a non-zero vector in $\mathcal{L}$ of length λ_1, we clearly have that $r > 0$. To now show that $\mathcal{L}$ is discrete it suffices to show that all sets are open, in particular for any $\mathbf{x} \in \mathcal{L}$, it suffices to show that $\{\mathbf{x}\}$ is open (since arbitrary unions of open sets are open). Taking the open ball of radius λ_1 around $\mathbf{x}$, it suffices to show that this ball uniquely intersects $\mathcal{L}$ in $\mathbf{x}$. If not, then there exists $\mathbf{y} \in \mathcal{L}$, $y \neq \mathbf{x}$ in $\mathcal{L}$ such that $\|\mathbf{x} - \mathbf{y}\| < r$. But then $\mathbf{x} - \mathbf{y}$ is a non-zero lattice vector of length less than $\lambda_1(\mathcal{L})$, a clear contradiction. □

The following exercise helps illustrate certain canonical situations where additive groups are or are not lattices.

Exercise 1

1. Let $A = \{x + \alpha y : x, y \in \mathbb{Z}\} \subseteq \mathbb{R}$, where $\alpha > 0$ is irrational. Show that A is not a lattice.
2. Let $\mathbf{v}_1, \ldots, \mathbf{v}_m \in \mathbb{Q}^n$. Show that $\mathcal{L}(\mathbf{v}_1, \ldots, \mathbf{v}_m)$ is a lattice (note that $\mathbf{v}_i$'s need not be linearly independent, in particular m maybe greater than n).

The following lemmas establish that non-trivial sets of type (Basis Rep.) or (Dual Rep.) admit lower bounds on λ_1 and hence are discrete. This corresponds to "half" the proof of

We note that the lower bounds presented below can vary greatly depending on the precise representation (there are in fact infinitely many such representations in general). One of the focus areas of the course will be to develop algorithms to produce "high quality" representations of a lattice, most often in the form of a "short" basis.

Lemma 4 *Let $\mathbf{B} \in \mathbb{R}^{k \times n}$, $k \geq 1$, be a non-singular matrix. Then $\lambda_1(\mathcal{L}(\mathbf{B})) \geq \sigma_{\min}(\mathbf{B}) > 0$, where $\sigma_{\min}(\mathbf{B})$ is the smallest singular value of $\mathbf{B}$.*

Proof: Recall that for any matrix $\mathbf{B}$, the minimum singular value is

$$\sigma_{\min}(\mathbf{B}) := \min_{\|\mathbf{x}\|=1} \|\mathbf{B}\mathbf{x}\| \ . \tag{1}$$

For $\mathbf{B}$ is non-singular, $\|\mathbf{B}\mathbf{x}\| \neq 0$ if $\mathbf{x} \neq \mathbf{0}$. Since $\|\mathbf{B}x\|$ is clearly continuous and the unit sphere in $\mathbb{R}^n$ is compact, the minimum in (1) is achieved, and hence $\sigma_{\min}(\mathbf{B}) > 0$ as needed.

Now let $\mathbf{x} \in \mathcal{L}(\mathbf{B})$ be a non-zero vector. We will show that $\|\mathbf{x}\| \geq \sigma_{\min}(\mathbf{B})$ to prove the lemma. Clearly, we may express $\mathbf{x} = \mathbf{Bz}$ where $\mathbf{z} \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$. Therefore

$$\|\mathbf{x}\| = \|\mathbf{Bz}\| \geq \sigma_{\min}(\mathbf{B})\|\mathbf{z}\| \geq \sigma_{\min}(\mathbf{B}) ,$$

where the last inequality follows since any non-zero integer vector has ℓ_2 norm at least 1 (since one of its coordinates has absolute value at least 1). □

Lemma 5 *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ with rows $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^n$ such that $\Lambda^\perp(\mathbf{A})$ is non-trivial. Then we have that $\lambda_1(\Lambda^\perp(\mathbf{A})) \geq \min_{i \in [m]} \frac{1}{\|\mathbf{a}_i\|} > 0$.*

Proof: Let $\mathbf{x} \in \Lambda^\perp(\mathbf{A})$ be a non-zero vector, which exists since by the non-triviality assumption. Since $\mathbf{x} \in \text{rowspan}(\mathbf{A})$, there exist $j \in [m]$ such that $\langle \mathbf{x}, \mathbf{a}_j \rangle \neq 0$. Since by definition $\langle \mathbf{x}, \mathbf{a}_j \rangle \in \mathbb{Z}$, we must in fact have $|\langle \mathbf{x}, \mathbf{a}_j \rangle| \geq 1$. Thus, by Cauchy-Schwarz

$$1 \leq |\langle \mathbf{x}, \mathbf{a}_j \rangle| \leq \|\mathbf{x}\| \|\mathbf{a}_j\| \Rightarrow \|\mathbf{x}\| \geq \frac{1}{\|\mathbf{a}_j\|} \geq \min_{i \in [m]} \frac{1}{\|\mathbf{a}_j\|} > 0 ,$$

as needed. □

5 Building a Lattice Basis

In this section, we will show that a lattice analogue of the basis extension theorem for linear subspaces holds. Recall that for a finite dimensional linear subspace, any subset of linearly independent vectors can be extended to a basis.

In the context of lattices, the direct analogue of this statement is unfortunately false. As an example, it is easy to see that the vector $2e_1$ cannot be extended to a basis of $\mathbb{Z}^2$, since the vector e_1 will have to be expressed as $\frac{1}{2}(2e_1)$ (which is not an integral combination) regardless of how we attempt to extend $2e_2$ to a basis. Note that the issue here is that the vector $2e_2$ can be scaled down to a lattice vector in $\mathbb{Z}^2$. Vectors for which this is not possible are called primitive. More precisely, $\mathbf{y} \in \mathcal{L}$ is said to be primitive for $\mathcal{L}$ if $\forall t \in \mathbb{R}$, $t\mathbf{y} \in \mathcal{L}$ if and only if $t \in \mathbb{Z}$. A natural question is thus whether any set of primitive lattice vectors can be extended to basis? Starting from exactly one vector, this will indeed turn out to be sufficient, as we will see below. However, already with two vectors the situation is not quite so simple. As another example, we may examine the vectors e_1 and $e_1 + 2e_2$. Both these vectors are primitive but do not generate $\mathbb{Z}^2$, since e_2 must be written as $\frac{1}{2}(e_1 + 2e_2) - \frac{1}{2}e_1$. To deal with this, we extend the notion of primitive to a set of vectors.

Definition 6 (Primitive) *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice. A set of linearly independent vectors $\mathbf{y}_1, \dots, \mathbf{y}_k \in \mathcal{L}$ is primitive w.r.t. $\mathcal{L}$ if $\mathcal{L} \cap \text{span}(\mathbf{y}_1, \dots, \mathbf{y}_k)$.*

The next theorem shows that any primitive set of vectors can indeed be extended to a lattice basis.

Theorem 7 *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a $k \geq 1$ dimensional lattice. Then $\mathcal{L}$ admits a basis of lattice vectors. Furthermore, given $\mathbf{b}_1, \dots, \mathbf{b}_i \in \mathcal{L}$ primitive w.r.t. to $\mathcal{L}$, there exists $\mathbf{b}_{i+1}, \dots, \mathbf{b}_k \in \mathcal{L}$ such that the extension $\mathbf{b}_1, \dots, \mathbf{b}_k$ is a basis of $\mathcal{L}$.*

Our strategy for proving the above will be to pick an arbitrary primitive vector in $\mathcal{L}$, continue inductively on the lattice projected orthogonal to this vector, and then lift the projected basis to the one for the full lattice. For this purpose we will need the following two lemmas. The first relates to when a projection of a lattice is a lattice, and the second explains how to lift a basis of a projected lattice.

LEMMA 8 *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice. Let $\mathbf{b}_1, \ldots, \mathbf{b}_k \in \mathcal{L}$ and $W := \mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_k)^\perp$. Then $\pi_W(\mathcal{L})$ is a lattice.*

PROOF: To show that $\pi_W(\mathcal{L})$ is a lattice we must show that it is an additive subgroup of $\mathbb{R}^n$ and that it is discrete. The additive subgroup property follows directly from linearity of π_W and that $\mathcal{L}$ is an additive subgroup. Thus, the main property to check is that $\pi_W(\mathcal{L})$ is discrete.

By Lemma 3, it suffices to show that for any bounded set $S \subseteq \mathbb{R}^n$ that $|S \cap \pi_W(\mathcal{L})| < \infty$. For this purpose, we will build an injective map $\tau : \pi_W(\mathcal{L}) \to \mathcal{L}$ satisfying that $\|\tau(\mathbf{x}) - \mathbf{x}\| \leq R$ where $R = \sum_{i=1}^k \frac{1}{2}\|\mathbf{b}_i\|$. Given such a τ, note that if $\mathbf{x} \in \pi_W(\mathcal{L}) \cap S$ then $\tau(\mathbf{x}) \in \mathcal{L} \cap (S + R\mathcal{B}_2^n)$. By injectivity of τ and discreteness of $\mathcal{L}$, we thus have that

$$|S \cap \pi_W(\mathcal{L})| \leq |(S + R\mathcal{B}_2^n) \cap \tau(\pi_W(\mathcal{L}))| \leq |(S + R\mathcal{B}_2^n) \cap \mathcal{L}| < \infty ,$$

as needed. Thus it suffices to build the required τ. To define τ on $\mathbf{x} \in \pi_W(\mathcal{L})$, we first pick $\hat{\mathbf{x}} \in \mathcal{L}$ such that $\pi_W(\hat{\mathbf{x}}) = \mathbf{x}$. Given that $\hat{\mathbf{x}} - \mathbf{x} \in W^\perp$ by construction, there exists a linear combination $\sum_{i=1}^k a_{\mathbf{x},i}\mathbf{b}_i = \hat{\mathbf{x}} - \mathbf{x}$ (since $\mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_k) = W^\perp$). We now define $\tau(\mathbf{x}) = \mathbf{x} + \sum_{i=1}^k (a_{\mathbf{x},i} - \lfloor a_{\mathbf{x},i} \rceil)\mathbf{b}_i$, where $\lfloor \cdot \rceil$ rounds to the nearest integer. We now show that τ satisfies the desired properties. That τ maps into $\mathcal{L}$ follows since

$$\tau(\mathbf{x}) = (\mathbf{x} + \sum_{i=1}^k a_{\mathbf{x},i}\mathbf{b}_i) - \sum_{i=1}^k \lfloor a_{\mathbf{x},i} \rceil \mathbf{b}_i = \underbrace{\hat{\mathbf{x}}}_{\in \mathcal{L}} - \sum_{i=1}^n \underbrace{\lfloor a_{\mathbf{x},i} \rceil \mathbf{b}_i}_{\in \mathcal{L}} \in \mathcal{L} .$$

Injectivity of τ follows immediately from the fact that $\pi_W(\tau(\mathbf{x})) = \mathbf{x}$, since $\mathbf{b}_1, \ldots, \mathbf{b}_k$ are in the kernel of π_W. Lastly, for the distance property, we see that

$$\|\tau(\mathbf{x}) - \mathbf{x}\| = \|\sum_{i=1}^k (a_{\mathbf{x},i} - \lfloor a_{\mathbf{x},i} \rceil)\mathbf{b}_i\| \leq \sum_{i=1}^k \frac{1}{2}\|\mathbf{b}_i\| = R ,$$

as needed. □

LEMMA 9 *Let $\mathcal{L} \subset \mathbb{R}^n$ be k-dimensional lattice. Assume that $\mathbf{b}_1, \ldots, \mathbf{b}_i \in \mathcal{L}$ is primitive w.r.t. $\mathcal{L}$ and that $\widetilde{\mathbf{b}}_{i+1}, \ldots, \widetilde{\mathbf{b}}_k \in \pi_W(\mathcal{L})$ is a basis of $\pi_W(\mathcal{L})$ where $W := \mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_i)^\perp$. Then for any choice of $\mathbf{b}_{i+1}, \ldots, \mathbf{b}_k \in \mathcal{L}$ such that $\pi_W(\mathbf{b}_j) = \widetilde{\mathbf{b}}_j$, $j \in \{i+1, \ldots, k\}$, the vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k$ form a basis of $\mathcal{L}$.*

PROOF: Take $\mathbf{x} \in \mathcal{L}$. To prove the lemma, we must show that we can express $\mathbf{x}$ as an integer combination of $\mathbf{b}_1, \ldots, \mathbf{b}_k$. Since these vectors are already linearly independent by construction, this will suffice to prove that they form a basis of $\mathcal{L}$.

Since by assumption $\widetilde{\mathbf{b}}_{i+1}, \ldots, \widetilde{\mathbf{b}}_k$ is a basis of $\pi_W(\mathcal{L})$, we may write $\pi_W(\mathbf{x}) = \sum_{j=i+1}^k z_j \widetilde{\mathbf{b}}_j$ where $z_{j+1}, \ldots, z_k \in \mathbb{Z}$. From here, we see that

$$\pi_W(\mathbf{x} - \sum_{j=i+1}^k z_j \mathbf{b}_j) = \pi_W(\mathbf{x}) - \sum_{j=i+1}^k z_j \widetilde{\mathbf{b}}_j = 0 \Rightarrow \mathbf{x} - \sum_{j=i+1}^k z_j \mathbf{b}_j \in \mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_i) .$$

In particular, $\mathbf{x} - \sum_{j=i+1}^{k} z_j \in \mathcal{L} \cap \operatorname{span}(\mathbf{b}_1, \ldots, \mathbf{b}_i)$, since $\mathcal{L}$ is an additive group. Since $\mathbf{b}_1, \ldots, \mathbf{b}_i$ is primitive w.r.t. $\mathcal{L}$, they form a basis for $\mathcal{L} \cap \operatorname{span}(\mathbf{b}_1, \ldots, \mathbf{b}_k)$, and hence we may write $\mathbf{x} - \sum_{j=i+1}^{k} z_j \mathbf{b}_j = \sum_{j=1}^{i} z_i \mathbf{b}_j$, for $z_1, \ldots, z_i \in \mathbb{Z}$. Rearranging, we get $\mathbf{x} = \sum_{j=1}^{k} z_j \mathbf{b}_j$, as needed. □

We are now ready to give a proof of the main theorem.

Proof of Theorem 7: We first show that $\mathcal{L}$ admits a basis. The proof will proceed by induction on the lattice dimension k.

To begin, we pick $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ arbitrarily (note that $\mathbf{y}$ exists since $\mathcal{L}$ has dimension at least 1). Let $\mathcal{L}_1 = \operatorname{span}(\mathbf{y}) \cap \mathcal{L}$. Clearly $\mathcal{L}_1$ is a 1 dimensional sublattice of $\mathcal{L}$. By Lemma 3, we can pick a $\mathbf{b}_1 \in \mathcal{L}_1 \setminus \{\mathbf{0}\}$ to be a shortest non-zero vector of $\mathcal{L}_1$. The next claim shows that $\mathbf{b}_1$ is primitive.

Claim 10 $\mathcal{L}_1 = \mathcal{L}(\mathbf{b}_1)$.

Proof: Assume not, then there exists $\mathbf{w} \in \mathcal{L}_1$ such that $\mathbf{w} \notin \mathbb{Z}\mathbf{b}_1$. Since $\mathcal{L}_1$ is 1-dimensional and $\mathbf{b}_1$ is non-zero, we have that $\mathbf{w} \in \operatorname{span}(\mathbf{b}_1)$ and hence $\mathbf{w} = \alpha \mathbf{b}_1$ for some $\alpha \in \mathbb{R} \setminus \mathbb{Z}$. Examine $\mathbf{w}' = (\alpha - \lfloor \alpha \rceil)\mathbf{b}_1$. Since $\lfloor \alpha \rceil \in \mathbb{Z}$ and $\mathbf{b}_1 \in \mathcal{L}_1$, we must have that $\mathbf{w}' = \mathbf{w} - \lfloor \alpha \rceil \mathbf{b}_1 \in \mathcal{L}_1$. Since $\alpha \in \mathbb{R} \setminus \mathbb{Z}$, we have that $\alpha - \lfloor \alpha \rceil \in (0, 1)$ and hence

$$\|\mathbf{b}_1\| > (\alpha - \lfloor \alpha \rceil)\|\mathbf{b}_1\| = \|\mathbf{w}'\| > 0,$$

a clear contradiction to $\mathbf{b}_1$ being a shortest non-zero vector of $\mathcal{L}_1$. Therefore $\mathcal{L}_1 = \mathbb{Z}\mathbf{b}_1 = \mathcal{L}(\mathbf{b}_1)$ as needed. □

Note that if $k = 1$, $\mathbf{b}_1$ is the desired basis, and so we are done. If $k \geq 2$, examine the projection $\pi_{\mathbf{b}_1^\perp}(\mathcal{L})$, which by Lemma 8 is indeed a lattice. Since $\mathcal{L}_2$ has dimension $k - 1$, by the induction hypothesis $\mathcal{L}_2$ admits a basis $\widetilde{\mathbf{b}}_2, \ldots, \widetilde{\mathbf{b}}_k$. By construction, we may choose $\mathbf{b}_2, \ldots, \mathbf{b}_k \in \mathcal{L}$ such that $\pi_{\mathbf{b}_1^\perp}(\mathbf{b}_j) = \widetilde{\mathbf{b}}_j \ \forall j \in \{2, \ldots, k\}$. By Lemma 9, $\mathbf{b}_1, \ldots, \mathbf{b}_k$ is the desired basis for $\mathcal{L}$, as needed.

Using the first part, we can now easly show that any primitive set of vector $\mathbf{b}_1, \ldots, \mathbf{b}_i \in \mathcal{L}$ can be extended to a basis. The proof is essentially identical to the proof in the previous paragraph, with the exception that we lift a basis from the projected lattice $\pi_{\operatorname{span}(\mathbf{b}_1, \ldots, \mathbf{b}_i)^\perp}(\mathcal{L})$. We leave the details to the reader. □

6 Proof of Theorem 2 (Lattice Representations)

We now have the tools to easily show that all definitions of lattices are equivalent. Note that if $\mathcal{L} = \{\mathbf{0}\}$, the statement is trivial, so we may assume that $\mathcal{L} \neq \{\mathbf{0}\}$.

Firstly, if $\mathcal{L} = \mathcal{L}(\mathbf{B})$ (i.e. Basis Rep. form), $\mathbf{B}$ non-singular, or $\mathcal{L} = \Lambda^\perp(\mathbf{A})$ (i.e. Dual Rep. form), we claim that $\mathcal{L}$ is indeed a lattice. In both cases, proving that $\mathcal{L}$ is an additive subgroup of $\mathbb{R}^n$ is direct, i.e. showing that $\forall \mathbf{x}, \mathbf{y} \in \mathcal{L}$, $-\mathbf{x} \in \mathcal{L}$, $\mathbf{x} + \mathbf{y} \in \mathcal{L}$, so we leave the details to the reader. To show discreteness, by Lemma 3 it suffices to show that $\lambda_1(\mathcal{L}) > 0$. For (Basis Rep.) a non-zero lower bound is given by Lemma 4, since B is non-singular, and for (Dual Rep.) by Lemma 5, since $\mathbf{A}$ must contains a least one non-zero row.

For the reverse direction, Theorem 7 shows that every lattice admits a basis. Thus, to finish the proof, it suffices to show that $\mathcal{L} = \mathcal{L}(\mathbf{B})$ can be expressed as $\Lambda^\perp(\mathbf{A})$. From basic linear algebra, $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_k)$ admits a corresponding dual basis $\mathbf{B}^* = (\mathbf{b}_1^*, \ldots, \mathbf{b}_k^*)$, satisfying $\langle \mathbf{b}_i, \mathbf{b}_j^* \rangle = 1$ if $i = j$ and 0 if $i \neq j$, and $\operatorname{span}(\mathbf{B}) = \operatorname{span}(\mathbf{B}^*)$. We now claim that $\mathcal{L}(\mathbf{B}) = \Lambda^\perp(\mathbf{B}^{*\mathsf{T}})$. Assume that

$\mathbf{x} \in \mathcal{L}(\mathbf{B})$, so $\mathbf{x} = \mathbf{B}\mathbf{z}$ for $\mathbf{z} \in \mathbb{Z}^k$. Then clearly $\mathbf{x} \in \mathrm{span}(\mathbf{B}) = \mathrm{span}(\mathbf{B}^*) = \mathrm{rowspan}(\mathbf{B}^{*\mathsf{T}})$. From here, we see that

$$\mathbf{B}^{*\mathsf{T}}\mathbf{x} = \mathbf{B}^{*\mathsf{T}}\mathbf{B}\mathbf{z} = \mathbf{I}_k\mathbf{z} = \mathbf{z} \in \mathbb{Z}^k \Rightarrow \mathbf{x} \in \Lambda^{\perp}(\mathbf{B}^{*\mathsf{T}}) \ ,$$

as needed. Now assume that $\mathbf{x} \in \Lambda^{\perp}(\mathbf{B}^{*\mathsf{T}})$. Letting $\mathbf{x}' = \mathbf{B}\mathbf{B}^{*\mathsf{T}}\mathbf{x}$, note that $\mathbf{x}' \in \mathcal{L}(\mathbf{B})$ since by assumption $\mathbf{B}^{*\mathsf{T}}\mathbf{x} \in \mathbb{Z}^k$. Thus it remains to show that $\mathbf{x} = \mathbf{x}'$. Since by construction $\mathbf{B}^{*\mathsf{T}}(\mathbf{x} - \mathbf{x}') = \mathbf{0}$ and both $\mathbf{x}, \mathbf{x}' \in \mathrm{span}(\mathbf{B}^*)$, we indeed get $\mathbf{x} = \mathbf{x}'$, as needed.

7 Equivalent Lattice Bases

From Theorem 7, we see that in fact a lattice can have many equivalent bases. A first question we ask is given two lattices bases $\mathbf{B}_1, \mathbf{B}_2$, when is it that $\mathcal{L}(\mathbf{B}_1) = \mathcal{L}(\mathbf{B}_2)$? In this section, we derive the basic relationship between equivalent lattices bases.

Definition 11 (Unimodular matrix) *A matrix* $\mathbf{U} \in \mathbb{Z}^{n\times n}$ *is unimodular if* $\det(\mathbf{U}) = \pm 1$.

For example, the matrix

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

is unimodular. The following lemma 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 12 $\mathbf{U} \in \mathbb{Z}^{n\times n}$ *is unimodular iff* $\mathbf{U}^{-1} \in \mathbb{Z}^{n\times n}$.

Proof: Assume $\mathbf{U}$ is unimodular. Let $\mathbf{M}_{ij} \in \mathbb{Z}^{(n-1)\times(n-1)}$ denote the principal minor of $\mathbf{U}$ obtained by deleting the i^{th} row and j^{th} column. Then by Cramer's rule, we know that $U^{-1}_{ij} = -1^{i+j}\det(\mathbf{M}_{ij})/\det(\mathbf{U})$. Since the determinant of an integer matrix is an integer, and since $\det(\mathbf{U}) = \pm 1$, we have that $\mathbf{U}^{-1} \in \mathbb{Z}^{n\times n}$ as needed.

Assume $\mathbf{U}^{-1} \in \mathbb{Z}^{n\times n}$. Then note that $1 = \det(\mathbf{I}) = \det(\mathbf{U}\mathbf{U}^{-1}) = \det(\mathbf{U})\det(\mathbf{U}^{-1})$. Since both $\mathbf{U}$ and $\mathbf{U}^{-1}$ are integer matrices, we know that both $\det(\mathbf{U}), \det(\mathbf{U}^{-1}) \in \mathbb{Z}$. Since the only integers dividing 1 are ± 1, we must have that $\det(\mathbf{U}) = \pm 1$ as needed. □

The next lemma tells us the two lattices bases generate the same lattice if and only if they are related by a unimodular transformation.

Lemma 13 *For non-singular matrices* $\mathbf{B}_1, \mathbf{B}_2 \in \mathbb{R}^{n\times k}$, $\mathcal{L}(\mathbf{B}_1) = \mathcal{L}(\mathbf{B}_2)$ *iff and only if* $\mathbf{B}_1 = \mathbf{B}_2\mathbf{U}$ *for some unimodular matrix* $\mathbf{U} \in \mathbb{Z}^{k\times k}$.

Proof: Assume $\mathcal{L}(\mathbf{B}_1) = \mathcal{L}(\mathbf{B}_2)$. Since each column of $\mathbf{B}_1$ is in $\mathcal{L}(\mathbf{B}_2)$, we can write $\mathbf{B}_1 = \mathbf{B}_2\mathbf{U}$ for some $\mathbf{U} \in \mathbb{Z}^{k\times k}$. Similarly, we get that $\mathbf{B}_1 = \mathbf{B}_2\mathbf{U}'$, for some $\mathbf{U}' \in \mathbb{Z}^{k\times k}$. Hence we get that $\mathbf{B}_1 = \mathbf{B}_2\mathbf{U} = \mathbf{B}_1\mathbf{U}'\mathbf{U}$. Since $\mathbf{B}_1$ is non-singular, $\mathbf{B}_1\mathbf{U}'\mathbf{U} = \mathbf{B}_1 \Leftrightarrow \mathbf{U}'\mathbf{U} = \mathbf{I}_k$, where $\mathbf{I}_k$ is the $k \times k$ identity. Hence $\mathbf{U}$ is unimodular as needed.

Assume $\mathbf{B}_1 = \mathbf{B}_2\mathbf{U}$ for some unimodular matrix $\mathbf{U} \in \mathbb{Z}^{k\times k}$. Clearly $\mathcal{L}(\mathbf{B}_1) \subseteq \mathcal{L}(\mathbf{B}_2)$. Since $\mathbf{U}$ is unimodular, $\mathbf{B}_2 = \mathbf{B}_1\mathbf{U}^{-1}$, where $\mathbf{U}^{-1} \in \mathbb{Z}^{k\times k}$, and hence $\mathcal{L}(\mathbf{B}_2) \subseteq \mathcal{L}(\mathbf{B}_1)$, as needed. □

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

The following lemma, which we give as an exercise, provides a different way to check whether two lattice bases are equivalent.

Exercise 2 Two bases are equivalent if and only if one can be obtained from the other by the following operations on columns:

1. $\mathbf{b}_i \leftarrow \mathbf{b}_i + k\mathbf{b}_j$ for some $k \in \mathbb{Z}$,

2. $\mathbf{b}_i \leftrightarrow \mathbf{b}_j$,

3. $\mathbf{b}_i \leftarrow -\mathbf{b}_i$.