COMPLEXITY OF LATTICE PROBLEMS

A Cryptographic Perspective

Daniele Micciancio
Shafi Goldwasser

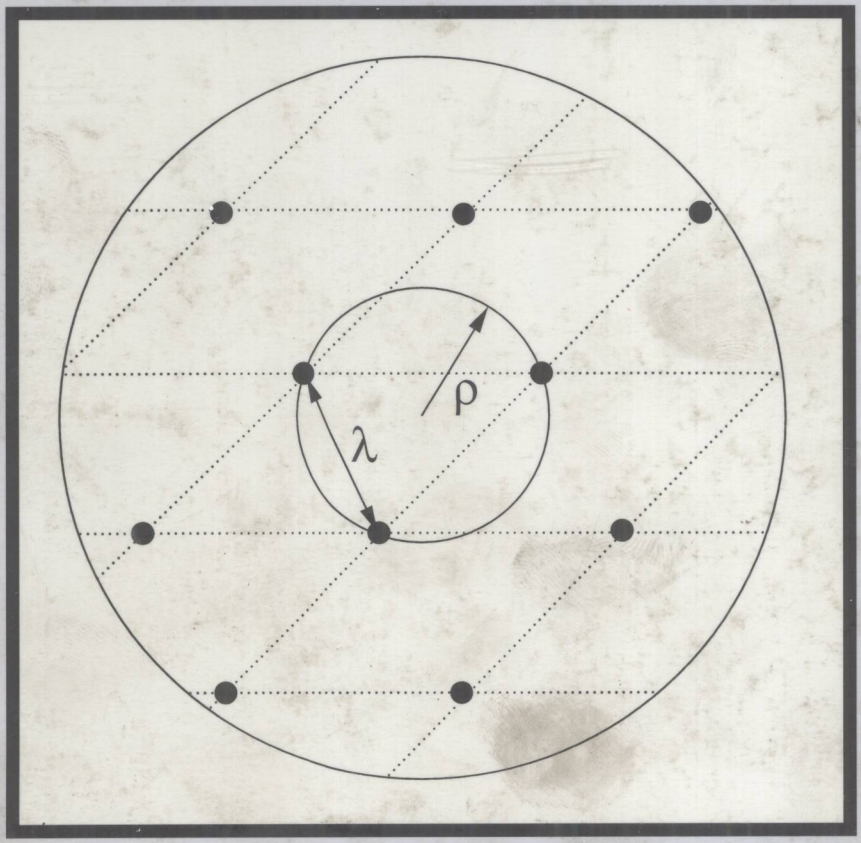

COMPLEXITY OF LATTICE PROBLEMS
A Cryptographic Perspective

Daniele Micciancio
University of California, San Diego

Shafi Goldwasser
The Massachusetts Institute of Technology, Cambridge

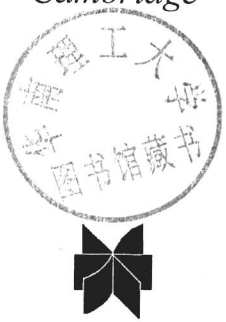

KLUWER ACADEMIC PUBLISHERS
Boston / Dordrecht / London

Distributors for North, Central and South America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA
Telephone (781) 871-6600
Fax (781) 681-9045
E-Mail < kluwer@wkap.com>

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS
Telephone 31 78 6392 392
Fax 31 78 6546 474
E-Mail < services@wkap.nl>

 Electronic Services < http://www.wkap.nl>

Library of Congress Cataloging-in-Publication Data

A C.I.P. Catalogue record for this book is available
from the Library of Congress.

Printed on acid-free paper.

Printed in Great Britain by IBT Global, London

COMPLEXITY OF LATTICE PROBLEMS
A Cryptographic Perspective

THE KLUWER INTERNATIONAL SERIES
IN ENGINEERING AND COMPUTER SCIENCE

Preface

Lattices are geometric objects that can be pictorially described as the set of intersection points of an infinite, regular n-dimensional grid. Despite their apparent simplicity, lattices hide a rich combinatorial structure, which has attracted the attention of great mathematicians over the last two centuries. Not surprisingly, lattices have found numerous applications in mathematics and computer science, ranging from number theory and Diophantine approximation, to combinatorial optimization and cryptography.

The study of lattices, specifically from a computational point of view, was marked by two major breakthroughs: the development of the LLL lattice reduction algorithm by Lenstra, Lenstra and Lovász in the early 80's, and Ajtai's discovery of a connection between the worst-case and average-case hardness of certain lattice problems in the late 90's.

The LLL algorithm, despite the relatively poor quality of the solution it gives in the worst case, allowed to devise polynomial time solutions to many classical problems in computer science. These include, solving integer programs in a fixed number of variables, factoring polynomials over the rationals, breaking knapsack based cryptosystems, and finding solutions to many other Diophantine and cryptanalysis problems.

Ajtai's discovery suggested a completely different way to use lattices in cryptography. Instead of using algorithmic solutions to computationally tractable lattice approximation problems to *break* cryptosystems, Ajtai's work shows how to use the existence of computationally intractable-to-approximate lattice problems to *build* cryptosystems which are *impossible to break*. Namely, design cryptographic functions that are provably as hard to break as it is to solve a computationally hard lattice problem.

Whereas in complexity theory we say that a problem is hard if it is hard for the worst case instance, in cryptography a problem is deemed hard only if it is hard in the average case (i.e., for all but a negligible

fraction of the instances). The novelty in Ajtai's result, is that he shows how to build a cryptographic function which is as hard to break on the average (e.g., over the random choices of the function instance) as it is to solve the worst case instance of a certain lattice problem. This achievement is unique to lattice theory at this time, and points to lattices as an ideal source of hardness for cryptographic purposes.

These new constructive applications of lattices, are deeply rooted in complexity theory, and were followed by a sharp increase in the study of lattices from a computational complexity point of view. This led to the resolution of several long standing open problems in the area. Most notably, the NP-hardness of the shortest vector problem in its exact and approximate versions. We present a self contained exposition of this latter result as well as other results on the computational complexity of lattice problems.

We did not attempt to cover everything known about lattices, as this would have filled several volumes. Rather, we selected a few representative topics, based on our personal taste and research experience. Regrettably, a topic which we neglect is duality and transference theorems. With this notable exception, we believe that most of the current ideas relevant to lattice based cryptography appear within in some form or another.

Many research questions regarding lattices and their cryptographic usage remain open. We hope that this book will help make lattice based cryptography more accessible to a wider audience, and ultimately yield further progress in this exciting research area.

Acknowledgments. Part of the material presented in this book is based on joint work of the authors with Shai Halevi, Oded Goldreich, Muli Safra and Jean-Pierre Seifert. Many other people have indirectly contributed to this book, either through their work, or through many conversations with the authors. Among them, we would like to mention Miklós Ajtai, Ravi Kannan, Amit Sahai, Claus Schnorr, Madhu Sudan and Salil Vadhan. We would like to thank all our coauthors and colleagues that have made this book possible.

The first author would like to thank also the National Science Foundation and Chris and Warren Hellman for partially supporting this work under NSF Career Award CCR-0093029 and a 2001-02 Hellman Fellowship.

DANIELE MICCIANCIO

Contents

Chapter 1

BASICS

This book is about algorithmic problems on point lattices, and their computational complexity. In this chapter we give some background about lattices and complexity theory.

1. Lattices

Let $\mathbb{R}^m$ be the m-dimensional Euclidean space. A *lattice* in $\mathbb{R}^m$ is the set

$$\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_n) = \left\{ \sum_{i=1}^{n} x_i \mathbf{b}_i : x_i \in \mathbb{Z} \right\} \tag{1.1}$$

of all integral combinations of n linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n$ in $\mathbb{R}^m$ $(m \geq n)$. The integers n and m are called the *rank* and *dimension* of the lattice, respectively. The sequence of vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n$ is called a *lattice basis* and it is conveniently represented as a matrix

$$\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_n] \in \mathbb{R}^{m \times n} \tag{1.2}$$

having the basis vectors as columns. Using matrix notation, (1.1) can be rewritten in a more compact form as

$$\mathcal{L}(\mathbf{B}) = \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^n\} \tag{1.3}$$

where $\mathbf{B}\mathbf{x}$ is the usual matrix-vector multiplication.

Graphically, a lattice can be described as the set of intersection points of an infinite, regular (but not necessarily orthogonal) n-dimensional grid. A 2-dimensional example is shown in Figure 1.1. There, the basis vectors are

$$\mathbf{b}_1 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \qquad \mathbf{b}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \tag{1.4}$$

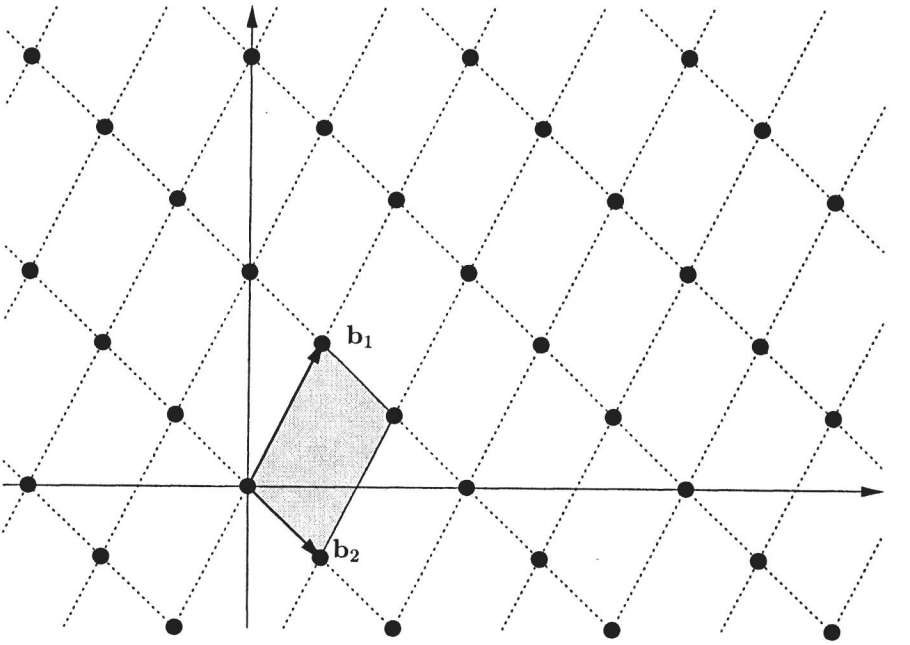

Figure 1.1. A lattice in $\mathbb{R}^2$

and they generate all the intersection points of the grid when combined with integer coefficients. The same lattice has many different bases. For example, vectors

$$\mathbf{b}'_1 = \mathbf{b}_1 + \mathbf{b}_2 = \begin{bmatrix} 2 \\ 1 \end{bmatrix}, \qquad \mathbf{b}'_2 = 2\mathbf{b}_1 + \mathbf{b}_2 = \begin{bmatrix} 3 \\ 3 \end{bmatrix} \tag{1.5}$$

are also a basis for lattice $\mathcal{L}(\mathbf{b}_1, \mathbf{b}_2)$. The grid generated by $\mathbf{b}'_1, \mathbf{b}'_2$ is shown in Figure 1.2. Notice that although the two grids are different, the set of intersection points is exactly the same, i.e., $\{\mathbf{b}_1, \mathbf{b}_2\}$ and $\{\mathbf{b}'_1, \mathbf{b}'_2\}$ are two different bases for the same lattice $\mathcal{L}(\mathbf{b}_1, \mathbf{b}_2) = \mathcal{L}(\mathbf{b}'_1, \mathbf{b}'_2)$.

Throughout the book, we use the convention that lattice points are always represented as *column* vectors. Wherever vectors are more conveniently written as rows, we use transpose notation. For example, the definition of vector $\mathbf{b}_1, \mathbf{b}_2$ in (1.4) can equivalently be rewritten as $\mathbf{b}_1 = [1, 2]^T, \mathbf{b}_2 = [1, -1]^T$, where $\mathbf{A}^T$ denotes the transpose of matrix $\mathbf{A}$.

A simple example of n-dimensional lattice is given by the set $\mathbb{Z}^n$ of all vectors with integral coordinates. A possible basis is given by the

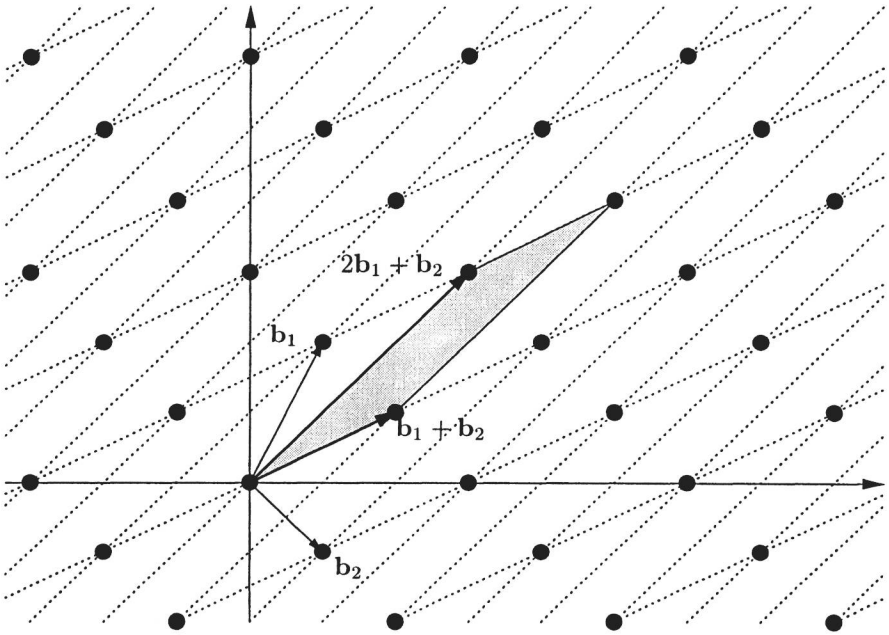

Figure 1.2. A different basis

standard unit vectors

$$\mathbf{e}_i = [\overbrace{0, \ldots, 0, \underbrace{1}_{i}, 0, \ldots, 0}^{n}]^T.$$

In matrix notation $\mathbb{Z}^n = \mathcal{L}(\mathbf{I})$ where $\mathbf{I} \in \mathbb{Z}^{n \times n}$ is the n-dimensional identity matrix, i.e., the $n \times n$ square matrix with 1's on the diagonal and 0's everywhere else.

When $n = m$, i.e., the number of basis vectors equals the number of coordinates, we say that $\mathcal{L}(\mathbf{B})$ is *full rank* or *full dimensional.* Equivalently, lattice $\mathcal{L}(\mathbf{B}) \subseteq \mathbb{R}^m$ is full rank if and only if the linear span of the basis vectors

$$\text{span}(\mathbf{B}) = \{\mathbf{Bx}: \mathbf{x} \in \mathbb{R}^n\} \tag{1.6}$$

equals the entire space $\mathbb{R}^m$. The difference between (1.3) and (1.6) is that while in (1.6) one can use arbitrary real coefficients to combine the basis vectors, in (1.3) only integer coefficients are allowed. It is easy to see that span($\mathbf{B}$) does not depend on the particular basis $\mathbf{B}$, i.e., if $\mathbf{B}$ and $\mathbf{B}'$ generate the same lattice then span($\mathbf{B}$) = span($\mathbf{B}'$). So,

for any lattice $\Lambda = \mathcal{L}(\mathbf{B})$, we can define the linear span of the lattice $\text{span}(\Lambda)$, without reference to any specific basis. Notice that $\mathbf{B}$ is a basis of $\text{span}(\mathbf{B})$ as a vector space. In particular, the rank of lattice $\mathcal{L}(\mathbf{B})$ equals the dimension of $\text{span}(\mathbf{B})$ as a vector space over $\mathbb{R}$ and it is a lattice invariant, i.e., it does not depend on the choice of the basis.

Clearly, any set of n linearly independent lattice vectors $\mathbf{B}' \in \mathcal{L}(\mathbf{B})$ is a basis for $\text{span}(\mathbf{B})$ as a vector space. However, $\mathbf{B}'$ is not necessarily a lattice basis for $\mathcal{L}(\mathbf{B})$. See Figure 1.3 for a 2-dimensional example. The picture shows the lattice $\mathcal{L}(\mathbf{b}_1, \mathbf{b}_2)$ generated by basis vectors (1.4) and the grid associated to lattice vectors

$$\mathbf{b}'_1 = \mathbf{b}_1 + \mathbf{b}_2 = \begin{bmatrix} 2 \\ 1 \end{bmatrix}, \qquad \mathbf{b}'_2 = \mathbf{b}_1 - \mathbf{b}_2 = \begin{bmatrix} 0 \\ 3 \end{bmatrix}. \tag{1.7}$$

Vectors $\mathbf{b}'_1$ and $\mathbf{b}'_2$ are linearly independent. Therefore, they are a basis for the plane $\mathbb{R}^2 = \text{span}(\mathbf{b}_1, \mathbf{b}_2)$ as a vector space. However, they are not a basis for $\mathcal{L}(\mathbf{b}_1, \mathbf{b}_2)$ because lattice point $\mathbf{b}_1$ cannot be expressed as an *integer* linear combination of $\mathbf{b}'_1$ and $\mathbf{b}'_2$. There is a simple geometric characterization for linearly independent lattice vectors that generate the whole lattice. For any n linearly independent lattice vectors $\mathbf{B}' = [\mathbf{b}'_1, \ldots, \mathbf{b}'_n]$ (with $\mathbf{b}'_i \in \mathcal{L}(\mathbf{B}) \subset \mathbb{R}^m$ for all $i = 1, \ldots, n$) define the half open parallelepiped

$$\mathcal{P}(\mathbf{B}') = \{\mathbf{B}'\mathbf{x} : 0 \leq x_i < 1\}. \tag{1.8}$$

Then, $\mathbf{B}'$ is a basis for lattice $\mathcal{L}(\mathbf{B})$ if and only if $\mathcal{P}(\mathbf{B}')$ does not contain any lattice vector other than the origin. Figures 1.1, 1.2 and 1.3 illustrate the two cases. The lattice in Figures 1.2 and 1.3 is the same as the one in Figure 1.1. In Figure 1.2, the (half open) parallelepiped $\mathcal{P}(\mathbf{B}')$ does not contain any lattice point other than the origin, and therefore $\mathcal{L}(\mathbf{B}') = \mathcal{L}(\mathbf{B})$. In Figure 1.3, parallelepiped $\mathcal{P}(\mathbf{B}')$ contains lattice point $\mathbf{b}_1$. Therefore $\mathcal{L}(\mathbf{B}') \neq \mathcal{L}(\mathbf{B})$ and $\mathbf{B}'$ is not a basis for $\mathcal{L}(\mathbf{B})$.

Notice that since $\mathbf{B}'$ is a set of linearly independent vectors, $\mathcal{L}(\mathbf{B}')$ is a lattice and $\mathbf{B}'$ is a basis for $\mathcal{L}(\mathbf{B}')$. Clearly, $\mathcal{L}(\mathbf{B}') \subseteq \mathcal{L}(\mathbf{B})$, i.e., any point from lattice $\mathcal{L}(\mathbf{B}')$ belongs also to lattice $\mathcal{L}(\mathbf{B})$. When $\mathcal{L}(\mathbf{B}') \subseteq \mathcal{L}(\mathbf{B})$, we say that $\mathcal{L}(\mathbf{B}')$ is a *sublattice* of $\mathcal{L}(\mathbf{B})$. If $\mathcal{L}(\mathbf{B}') = \mathcal{L}(\mathbf{B})$ we say that bases $\mathbf{B}$ and $\mathbf{B}'$ are *equivalent*. If $\mathcal{L}(\mathbf{B}') \subseteq \mathcal{L}(\mathbf{B})$, but $\mathcal{L}(\mathbf{B}') \neq \mathcal{L}(\mathbf{B})$, then bases $\mathbf{B}$ and $\mathbf{B}'$ are not equivalent, and $\mathcal{L}(\mathbf{B}')$ is a *proper* sublattice of $\mathcal{L}(\mathbf{B})$.

Equivalent bases (i.e., bases that generate the same lattice) can be algebraically characterized as follows. Two bases $\mathbf{B}, \mathbf{B}' \in \mathbb{R}^{m \times n}$ are equivalent if and only if there exists a unimodular matrix $\mathbf{U} \in \mathbb{Z}^{n \times n}$ (i.e., an integral matrix with determinant $\det(\mathbf{U}) = \pm 1$) such that $\mathbf{B}' = \mathbf{B}\mathbf{U}$. The simple proof is left to the reader as an exercise.

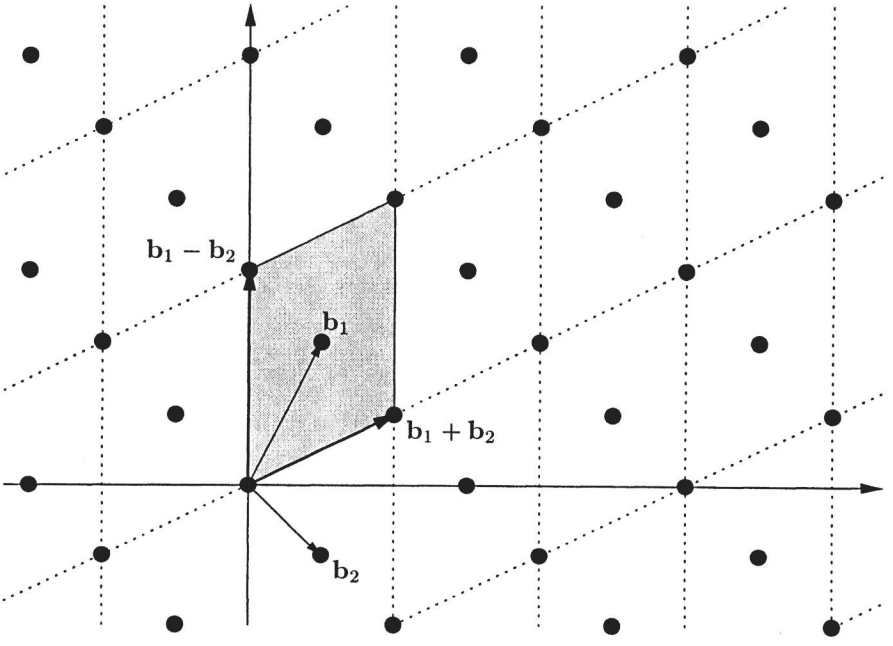

Figure 1.3. The sublattice generated by $\mathbf{b}_1 + \mathbf{b}_2$ and $\mathbf{b}_1 - \mathbf{b}_2$

When studying lattices from a computational point of view, it is customary to assume that the basis vectors (and therefore any lattice vector) have all rational coordinates. It is easy to see that rational lattices can be converted to integer lattices (i.e., sublattices of $\mathbb{Z}^n$) by multiplying all coordinates by an appropriate integer scaling factor. So, without loss of generality, in the rest of this book we concentrate on integer lattices, and, unless explicitly stated otherwise, we always assume that lattices are represented by a basis, i.e., a matrix with integer coordinates such that the columns are linearly independent.

Lattices can also be characterized without reference to any basis. A lattice can be defined as a discrete nonempty subset Λ of $\mathbb{R}^m$ which is closed under subtraction, i.e., if $\mathbf{x} \in \Lambda$ and $\mathbf{y} \in \Lambda$, then also $\mathbf{x} - \mathbf{y} \in \Lambda$. Here "discrete" means that there exists a positive real $\lambda > 0$ such that the distance between any two lattice vectors is at least λ. A typical example is the set $\Lambda = \{\mathbf{x} \in \mathbb{Z}^n : \mathbf{A}\mathbf{x} = \mathbf{0}\}$ of integer solutions of a system of homogeneous linear equations. Notice that Λ always contains the origin $\mathbf{0} = \mathbf{x} - \mathbf{x}$, it is closed under negation (i.e., if $\mathbf{x} \in \Lambda$ then $-\mathbf{x} = \mathbf{0} - \mathbf{x} \in \Lambda$), and addition (i.e., if $\mathbf{x}, \mathbf{y} \in \Lambda$ then $\mathbf{x} + \mathbf{y} = \mathbf{x} - (-\mathbf{y}) \in \Lambda$). In other words, Λ is a discrete additive subgroup of $\mathbb{R}^m$.

1.1 Determinant

The *determinant* of a lattice $\Lambda = \mathcal{L}(\mathbf{B})$, denoted $\det(\Lambda)$, is the n-dimensional volume of the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$ spanned by the basis vectors. (See shaded areas in Figures 1.1 and 1.2.) The determinant is a lattice invariant, i.e., it does not depend on the particular basis used to compute it. This immediately follows from the characterization of equivalent bases as matrices $\mathbf{B}' = \mathbf{BU}$ related by a unimodular transformation $\mathbf{U}$. Geometrically, this corresponds to the intuition that the (n-dimensional) volume of the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$ equals the inverse of the density of the lattice points in $\text{span}(\mathbf{B})$. As an example consider the bases in Figures 1.1 and 1.2. The areas of the fundamental regions (i.e., the shaded parallelepipeds in the pictures) are exactly the same because the two bases generate the same lattice. However, the shaded parallelepiped in Figure 1.3 has a different area (namely, twice as much as the original lattice) because vectors (1.7) only generate a sublattice.

A possible way to compute the determinant is given by the usual *Gram-Schmidt orthogonalization* process. For any sequence of vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n$, define the corresponding Gram-Schmidt orthogonalized vectors $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ by

$$\mathbf{b}_i^* = \mathbf{b}_i - \sum_{j=1}^{i-1} \mu_{i,j} \mathbf{b}_j^* \tag{1.9a}$$

$$\mu_{i,j} = \frac{\langle \mathbf{b}_i, \mathbf{b}_j^* \rangle}{\langle \mathbf{b}_j^*, \mathbf{b}_j^* \rangle} \tag{1.9b}$$

where $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{m} x_i y_i$ is the inner product in $\mathbb{R}^m$. For every i, $\mathbf{b}_i^*$ is the component of $\mathbf{b}_i$ orthogonal to $\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}$. In particular, $\text{span}(\mathbf{b}_1, \ldots, \mathbf{b}_i) = \text{span}(\mathbf{b}_1^*, \ldots, \mathbf{b}_i^*)$ and vectors $\mathbf{b}_i^*$ are pairwise orthogonal, i.e., $\langle \mathbf{b}_i^*, \mathbf{b}_j^* \rangle = 0$ for all $i \neq j$. The determinant of the lattice equals the product of the lengths of the orthogonalized vectors

$$\det(\mathcal{L}(\mathbf{B})) = \prod_{i=1}^{n} \|\mathbf{b}_i^*\| \tag{1.10}$$

where $\|\mathbf{x}\| = \sqrt{\sum_i x_i^2}$ is the usual Euclidean length. We remark that the definition of the orthogonalized vectors $\mathbf{b}_i^*$ depends on the order of the original basis vectors. Given basis matrix $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$, we denote by $\mathbf{B}^*$ the matrix whose columns are the orthogonalized vectors $[\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*]$. Clearly, $\mathbf{B}^*$ is a basis of $\text{span}(\mathbf{B})$ as a vector space. However, $\mathbf{B}^*$ is not usually a lattice basis for $\mathcal{L}(\mathbf{B})$. In particular, not every lattice has a basis consisting of mutually orthogonal vectors.

Notice that if the $\mathbf{b}_i$'s are rational vectors (i.e., vectors with rational coordinates), then also the orthogonalized vectors $\mathbf{b}_i^*$ are rationals. If lattice $\mathcal{L}(\mathbf{B})$ is full dimensional (i.e. $m = n$), then $\mathbf{B}$ is a nonsingular square matrix and $\det(\mathcal{L}(\mathbf{B}))$ equals the absolute value of the determinant of the basis matrix $\det(\mathbf{B})$. For integer lattices, $\mathbf{B}$ is a square integer matrix, and the lattice determinant $\det(\mathcal{L}(\mathbf{B})) = \det(\mathbf{B})$ is an integer. In general, the reader can easily verify that $\det(\mathcal{L}(\mathbf{B}))$ equals the square root of the determinant of the Gram matrix $\mathbf{B}^T\mathbf{B}$, i.e., the $n \times n$ matrix whose (i,j)th entry is the inner product $\langle \mathbf{b}_i, \mathbf{b}_j \rangle$:

$$\det(\mathcal{L}(\mathbf{B})) = \sqrt{\det(\mathbf{B}^T\mathbf{B})}. \tag{1.11}$$

This gives an alternative way to compute the determinant of a lattice (other than computing the Gram-Schmidt orthogonalized vectors), and shows that if $\mathbf{B}$ is an integer matrix, then the determinant of $\mathcal{L}(\mathbf{B})$ is always the square root of a positive integer, even if $\det(\mathcal{L}(\mathbf{B}))$ is not necessarily an integer when the lattice is not full rank.

1.2 Successive minima

Let $\mathcal{B}_m(\mathbf{0}, r) = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\| < r\}$ be the m-dimensional open ball of radius r centered in $\mathbf{0}$. When the dimension m is clear from the context, we omit the subscript m and simply write $\mathcal{B}(\mathbf{0}, r)$. Fundamental constants associated to any rank n lattice Λ are its successive minima $\lambda_1, \ldots, \lambda_n$. The ith minimum $\lambda_i(\Lambda)$ is the radius of the smallest sphere centered in the origin containing i linearly independent lattice vectors

$$\lambda_i(\Lambda) = \inf\{r : \dim(\operatorname{span}(\Lambda \cap \mathcal{B}(\mathbf{0}, r))) \geq i\}. \tag{1.12}$$

Successive minima can be defined with respect to any norm. A norm is a positive definite, homogeneous function that satisfies the triangle inequality, i.e., a function $\|\cdot\| : \mathbb{R}^n \to \mathbb{R}$ such that

- $\|\mathbf{x}\| \geq 0$ with equality only if $\mathbf{x} = \mathbf{0}$
- $\|\alpha\mathbf{x}\| = |\alpha| \cdot \|\mathbf{x}\|$
- $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$. An important family of norm functions is given by the ℓ_p norms. For any $p \geq 1$, the ℓ_p norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^{n} x_i^p \right)^{1/p}. \tag{1.13a}$$