

INTEGER PROGRAMMING, LATTICE ALGORITHMS,
AND DETERMINISTIC VOLUME ESTIMATION

A Dissertation
Presented to
The Academic Faculty

by

Daniel Nicolas Dadush

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in
Algorithms, Combinatorics, and Optimization

H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
August 2012

INTEGER PROGRAMMING, LATTICE ALGORITHMS, AND DETERMINISTIC VOLUME ESTIMATION

Approved by:

Dr. Santosh S. Vempala, Advisor
School of Computer Science
Georgia Institute of Technology

Dr. Christopher Peikert
School of Computer Science
Georgia Institute of Technology

Dr. Santanu Dey
H. Milton Stewart School of Industrial
and Systems Engineering
Georgia Institute of Technology

Dr. Arkadi Nemirovski
H. Milton Stewart School of Industrial
and Systems Engineering
Georgia Institute of Technology

Dr. William Cook
H. Milton Stewart School of Industrial
and Systems Engineering
Georgia Institute of Technology

Dr. Daniele Micciancio
School of Computer Science
University of California, San Diego

Date Approved: June 14th, 2012

*Pour ma famille, et ma mamie,
qui ont toujours eu confiance en moi.*

ACKNOWLEDGEMENTS

Every PhD is a journey, and I have been fortunate to have had the support and guidance of many throughout my journey at Georgia Tech. First and foremost, I would like to thank my advisor, Santosh Vempala, for introducing me to the beautiful world of convex geometry. It is through his generous support and guidance that I learned how to tackle big problems and set an ambitious research agenda for myself, a set of skills that will serve me well for the rest of my career. I would also like to deeply thank Chris Peikert for teaching me about lattices and their mysterious properties, for being an amazing collaborator (especially on those long days before submission), and whose advocacy and guidance has been invaluable to me. Santanu Dey and Juan Pablo Vielma, for countless conversations about life, research, and all that jazz, and for being great friends and collaborators. Bill Cook, for agreeing to be my reader, and for being the first (and probably last) person to plug one of my talks on Twitter. I would like to thank all the members of committee, consisting of my advisor, Chris Peikert, Daniele Micciancio, Santanu Dey, Bill Cook and Arkadi Nemirovski.

I would like to thank all those who encouraged me to pursue my graduate studies. In particular, my thanks go out to Eli Upfal and Pascal Van Hentenryck, who were my mentors while I was at Brown University. I would also like to thank George Nemhauser, for inviting me to visit Georgia Tech, and allowing me to learn about the ACO program.

I am extremely grateful for the support I've received from both the ACO program and my home department ISyE. My thanks go out to Robin Thomas and Gary Parker, for making this institutional support available to me, which was so much more than I could have ever asked for. I will have to work hard during the years to come to

make up my debt to them and to the ACO program as a whole. I would like to thank the many administrative assistants who have made managing the bureaucracy and organizing events at Tech so much easier: Pam Morrison, Anita Race, Elizabeth Ndongi, Dani Denton, Annette Rohrs, Inetta Worthy, and Sharon McDowell.

I have greatly benefitted from research collaborations and discussions with many people throughout the years. I would like to thank Arkadi Nemirovski, for always patiently entertaining my fanciful ideas, and answering my questions about convex analysis with remarkable speed and detail (occasionally as multiple page pdfs!). Daniele Micciancio, whose ingenuity and speed with complex lattice structures has left me amazed, and with whom I'm hopeful for many fruitful collaborations in the future. Gideon Schechtman, for hosting me at the Weizmann Institute and for letting me know that cubes tile better than balls. Boaz Klartag, for being a generous mentor while visiting Tel Aviv University, opening up my eyes to whole new areas in convex geometry, and suggesting that I make his M-Ellipsoid construction algorithmic. Matthias Köppe, for hosting me at UC Davis and introducing me to reverse search. Grigoris Paouris, for hosting me at Texas A&M and teaching me about the ℓ -Ellipsoid. Gabor Kun, for solving an open problem of mine in less than five minutes and for a great bottle of Tokaji wine. Friedrich Eisenbrand, for hosting me at EPFL and getting me hooked onto discrepancy theory. Among the many others I would like to thank are Karthik Chandrasekaran, Aditya Bhaskara, Ravishankar Krishnaswamy, Kunal Talwar, Moritz Hardt, Marco Molinaro, Aravindan Vijayaraghavan, Nikhil Bansal, Maxim Sviridenko, Konstantin Makarychev, Sebastian Pokutta, Oded Regev, Assaf Naor, Andreas Dress, and many more.

I would also like to thank the many friends I've had while at Tech, who helped make my life in graduate school incredibly fun, refreshing, and stimulating: Luke Postle, Megan Hodge, Arash Asadi, Noah & Amanda Streib, Karthik Chandrasekaran, Linji & Pengyi Yang, Xue Feng, László & Ágota Véghe, Lev Reysin, Elena Grigorescu,

Alejandro Toriello, Fatma Kilinc-Karzan, Dimitri Papageorgiou, Steve Tyber, Kael Stilp, and so many more.

To my wonderful family, my parents Uri and Gilda and my sister Sarah, I would like to say thank you from the bottom of my heart for being my bedrock of love and support, without which I could not have come this far. Lastly, I send my love to my girlfriend, Kristin Webb, for giving me a beautiful future to look forward to together in the years to come.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF FIGURES	xi
SUMMARY	xii
I INTRODUCTION	1
1.1 Contributions of this Thesis	2
1.1.1 Cutting Planes	2
1.1.2 Algorithmic Convex Geometry	3
1.1.3 Lattice Problems	5
1.1.4 Integer Programming	6
II MATHEMATICAL BACKGROUND	9
2.1 Basics	9
2.1.1 Linear Spaces	12
2.1.2 Elementary Analysis	18
2.1.3 Probability and Measure	19
2.2 Convexity	25
2.3 Convex Geometry	30
2.3.1 Logconcave Functions and Convex Bodies	31
2.3.2 Geometric Inequalities	32
2.4 Lattices	35
2.4.1 Packing, Covering and Tiling	38
2.4.2 Lattice Inequalities	41
2.5 Computational Complexity	42
2.5.1 Computational Model	44
2.5.2 Operations on Convex Bodies	47
2.5.3 Fundamental Algorithms	48

III ON THE CHVÁTAL-GOMORY CLOSURE OF A COMPACT CONVEX SET	50
3.1 Introduction	50
3.2 Definitions, Main Result and Proof Idea	52
3.3 $CG(K, S) \cap H_{\mathbf{v}} = CG(F_{\mathbf{v}})$ and $CG(K, S) \subseteq H_{\mathbf{v}}^{\leq}$	54
3.3.1 Lifting CG Cuts	54
3.3.2 Separating All Points in $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$	59
3.3.3 Lifting the CG Closure of an Exposed Face of K	69
3.4 Approximation of the CG Closure	70
3.4.1 Approximation 1 of the CG Closure	70
3.4.2 Approximation 2 of the CG Closure	73
3.5 Proof of Theorem	74
3.6 Conclusion	77
IV THE M-ELLIPSOID AND VOLUME ESTIMATION	79
4.1 Introduction	79
4.1.1 Results	81
4.1.2 The M-Ellipsoid	85
4.2 Building a Covering	87
4.3 Klartag's Construction	96
4.3.1 A Las Vegas Algorithm for Generating an M-Ellipsoid	102
4.3.2 Helper Algorithms	105
4.3.3 Geometric Estimates	109
4.4 Milman's Construction	114
4.4.1 The Lewis Ellipsoid	116
4.4.2 Covering Numbers and Volume Estimates	118
4.4.3 A Deterministic M-Ellipsoid Construction	119
4.4.4 Analysis	123
4.5 An Asymptotically Optimal Volume Algorithm	130
4.6 Computing an Approximate Center of Mass	134

4.7	Conclusion	137
V	EFFICIENT DETERMINISTIC ALGORITHMS FOR LATTICE PROBLEMS	139
5.1	Introduction	139
5.1.1	Lattice Problems	141
5.1.2	Results and Techniques	142
5.2	Lattice Point Enumeration in Convex Bodies	148
5.3	Shortest Vector Problem	164
5.4	Closest Vector Problem	168
5.5	Approximate Closest Vector Problem	172
5.5.1	A Simple Randomized Lattice Sparsifier Construction	177
5.5.2	Derandomizing the Lattice Sparsifier Construction	182
5.6	Finding a Central Lattice Point	192
5.7	Conclusion	193
VI	GEOMETRY OF THE DISCRETE GAUSSIAN AND THE FLATNESS THEOREM	196
6.1	Introduction	196
6.2	Preliminaries	199
6.3	Bounding the Smoothing Parameter	201
6.4	Comparing the Discrete and Continuous Gaussian	207
6.5	Flatness Theorem Proof	212
6.6	Conclusion	214
VII	THE INTEGER PROGRAMMING PROBLEM	216
7.1	Introduction	216
7.1.1	Results	218
7.2	Preprocessing the Integer Program	223
7.3	An Improved Lenstra Type Algorithm	227
7.4	An Improved Kannan Type Algorithm	231
7.4.1	Finding a Thin Projection	232

7.4.2	The Improved Algorithm	240
7.5	Convex Integer Minimization	245
7.6	Conclusion	255
	REFERENCES	258
	VITA	267

LIST OF FIGURES

4.1	The M-Ellipsoid Algorithm	123
4.2	Deterministic Volume Algorithm	132

SUMMARY

The main subject of this thesis is the development of new geometric tools and techniques for solving classic problems within the geometry of numbers and convex geometry. At a high level, the problems considered in this thesis concern the varied interplay between the continuous and the discrete, an important theme within computer science and operations research.

The first subject we consider is the study of cutting planes for non-linear integer programs. Cutting planes have been implemented to great effect for linear integer programs, and so understanding their properties in more general settings is an important subject of study. As our contribution to this area, we show that Chvátal-Gomory closure of any compact convex set is a rational polytope. As a consequence, we resolve an open problem of Schrijver (Ann. Disc. Math. '80) regarding the same question for irrational polytopes.

The second subject of study is that of ellipsoidal approximation of convex bodies. Different such notions have been important to the development of fundamental geometric algorithms: e.g. the ellipsoid method for convex optimization (enclosing ellipsoids), or random walk methods for volume estimation (inertial ellipsoids). Here we consider the construction of an ellipsoid with good “covering” properties with respect to a convex body, known in convex geometry as the M -ellipsoid. As our contribution, we give two algorithms for constructing M -ellipsoids, and provide an application to near-optimal deterministic volume estimation in the oracle model.

Equipped with this new geometric tool, we move to the study of classic lattice problems in the geometry of numbers, namely the Shortest (SVP) and Closest Vector Problems (CVP). Here we use M -ellipsoid coverings, combined with an algorithm of

Micciancio and Voulgaris for CVP in the ℓ_2 norm (STOC '10), to obtain the first deterministic $2^{O(n)}$ time algorithm for the SVP in general norms. Combining this algorithm with a novel lattice sparsification technique, we derive the first deterministic $2^{O(n)}(1 + 1/\epsilon)^n$ time algorithm for $(1 + \epsilon)$ -approximate CVP in general norms.

For the next subject of study, we analyze the geometry of general integer programs. A central structural result in this area is Kinchine's flatness theorem, which states that every lattice free convex body has integer width bounded by a function of dimension. As our contribution, we build on the work Banaszczyk, using tools from lattice based cryptography, to give a new and tighter proof of the flatness theorem.

Lastly, combining all the above techniques, we consider the study of algorithms for the Integer Programming Problem (IP). As our main contribution, we give a new $2^{O(n)}n^n$ time algorithm for IP, which yields the fastest currently known algorithm for IP and improves on the classic works of Lenstra (MOR '83) and Kannan (MOR '87).

CHAPTER I

INTRODUCTION

Throughout the twentieth century, the study of convexity and its interactions with discrete structures, from both the algorithmic and structural viewpoint, has led to many fundamental discoveries in mathematics. Major achievements range from the development of the geometry of numbers by Minkowski, to the development of efficient algorithms for optimization problems within operations research and computer science.

In this thesis, we will be concerned with problems involving the intersections of continuous and discrete spaces. In particular, we will examine algorithmic problems over lattices (discrete subgroups of \mathbb{R}^n) and convex bodies. Our main focus problems will be on lattice problems, namely, the Shortest (SVP) and Closest Vector Problems (CVP), as well as the Integer Programming Problem (IP). The combined algorithmic study of these problems was coined the algorithmic geometry of numbers by Kannan [72]. In Kannan's view, this subject, which combined the classic mathematical theory of Minkowski [97], as well as the sophisticated and elegant algorithmics of researchers such as Lenstra [84], and Lovász [82], and certainly himself [70], was perhaps the perfect playground for testing new combinatorial and geometric techniques on exceedingly difficult computational problems.

Twenty five after Kannan has written his survey on the subject, there has been tremendous progress on both sides of the equation: the subject of lattice algorithms has grown by leaps and bounds due to its many new connections to cryptography and complexity, and our understanding of the asymptotic properties of convex bodies (volume, concentration, isoperimetry, etc.) has grown beyond measure. However,

the essential synthesis for which Kannan had hoped, i.e. of bringing these new understandings together in a singular and focused manner to bear on the core set of problems within the subject, remains undone.

The main goal of this thesis, is to take a step in this direction. The theme of combining newly developed geometry, with recent techniques on lattices, will be pervasive throughout. In the chapters that follow, we will revisit nearly every classical problem within the subject from a fresh perspective, and hope to point the way towards future progress.

1.1 Contributions of this Thesis

1.1.1 Cutting Planes

The study of cutting planes, i.e. valid linear inequalities for the feasible solutions to an integer program, has played a fundamental role in the development of IP technology. The main body of cutting plane research has, until recently, focused on integer linear programming models (ILP). Due to the great success of cutting planes for linear models, and the increasing need to solve non-linear IPs, generalizing the theory of cuttings planes to more general settings is an important task.

In Chapter 3, with this aim in mind, we examine the properties of cutting planes for convex integer programs, i.e. IPs where the feasible region of the continuous relaxation is a general convex set. One of the first and most important classes of cutting planes developed for IP are the Chvátal-Gomory (CG) cuts, which were first introduced in [55] to design the first cutting plane algorithm for ILP.

The basis of our study is a fundamental structural result due to Schrijver [119], which states that the CG closure of a rational polyhedron, i.e. the set obtained by adding all CG cuts over the feasible region, is a rational polyhedron. A natural question is whether Schrijver's result extends to more general settings, and indeed Schrijver himself poses the question of whether it holds for polytopes described by

irrational data. As our main contribution, we extend Schrijver’s result, and show that the CG closure of any compact convex set is a rational polytope. This resolves Schrijver’s original question, and helps extend the theory of cutting planes to the convex IP setting.

This Chapter is based on joint work with Santanu Dey and Juan Pablo Vielma which appeared in the proceeding of the conference on *Integer Programming and Combinatorial Optimization*, 2011 [34].

1.1.2 Algorithmic Convex Geometry

In Chapter 4, we develop a nearly optimal algorithm for deterministically estimating the volume of a convex body. To achieve this, we develop new geometric algorithms which may find wider application, and which are used throughout the rest of thesis. At a high level, our main contributions are to make algorithmic certain fundamental constructions from convex geometry and the local theory of Banach spaces.

The M-Ellipsoid. An n dimensional ellipsoid is a body representable as an affine transformation of the euclidean ball (see Section 2.2 for precise definitions). Different types of ellipsoidal approximations for convex bodies, and techniques to compute them, have been fundamental to the development of algorithms in Operations Research and Computer Science. For example, the ellipsoid method for convex optimization, which was used to give the first polynomial algorithm for Linear Programming [73], is based on constructing a shrinking sequence of enclosing ellipsoids for the feasible region.

As a first problem, we examine the problem of building an M-Ellipsoid for a n dimensional convex body K . An M-Ellipsoid E for K approximates K from a “covering” perspective. Precisely, E is an M-Ellipsoid for K if $2^{O(n)}$ translates of E suffices to cover K and vice versa. The existence of the M-Ellipsoid was first proved by Milman [92], who used it to derive many fundamental inequalities in convex geometry.

Our motivation for constructing the M-Ellipsoid will be due to the utility of its implied covering, for which we give applications to volume estimation, lattice algorithms, and integer programming.

Our first results, are to give two different algorithms for computing M-Ellipsoids for convex bodies presented by membership oracles (i.e. where we can query whether a point is inside the body; see Section 2.5.1 for precise definitions). Our first algorithm, based on a construction of Klartag [77], runs in randomized polynomial time and succeeds with high probability. Our second algorithm, based on Milman’s original construction [92], is deterministic and runs in $2^{O(n)}$ time and uses polynomial space. Furthermore, we show that there is a $2^{\Omega(n)}$ lower bound for any deterministic construction in the oracle model.

Our next result within this context is a near optimal algorithm for computing ellipsoid coverings. Given as input an n dimensional convex body K and ellipsoid E , we give a polynomial space algorithm for outputting the translates of E corresponding to a covering of K by E . Furthermore, the size of the covering outputted by the algorithm is at most a $2^{O(n)}$ factor larger than the optimal such covering, and the runtime of the algorithm is proportional to the size of the outputted covering.

Deterministic Volume Estimation. The problem of estimating the volume of a convex body is a central problem in computer science, which has lead to the development of many algorithmic techniques for studying high dimensional distributions. The volume estimation problem is one of the prime examples for the power of randomization. In the work of Elekes [48], and later Bárány and Füredi [51, 52], it was shown that deterministically computing volume in the oracle model is *hard*. In particular, any deterministic algorithm which reliably estimates the volume of symmetric convex bodies to within a factor $(1 + \epsilon)^n$ must make at least $(1 + \frac{1}{\epsilon})^{\Omega(n)}$ membership queries in the worst case. In contrast, the breakthrough result of Dyer, Frieze, and

Kannan [45] gives a $\text{poly}(n, \frac{1}{\epsilon})$ time randomized algorithm which estimates volume to within $(1 + \epsilon)$ with high probability using Monte Carlo markov chain techniques.

As our contribution in this area, we show that the lower bounds for volume estimation are essentially tight for all values of $0 < \epsilon \leq 1$. Precisely, we use a variant of the deterministic M-Ellipsoid construction, together with the ellipsoid covering algorithm, to give a deterministic $(1 + \frac{1}{\epsilon})^{O(n)}$ time and polynomial space algorithm for estimating the volume of any symmetric convex body to within a factor $(1 + \epsilon)^n$.

The results in this Chapter are based two works, the first is joint with Chris Peikert and Santosh Vempala and appeared in the *Symposium on Foundations of Computer Science*, 2011 [36], and the second is joint with Santosh Vempala [32].

1.1.3 Lattice Problems

The Shortest (SVP) and Closest Vector Problems (CVP) are classic problems in the geometry of numbers and computer science. Given a n dimensional lattice \mathcal{L} (a lattice is a discrete subgroup of \mathbb{R}^n , see section 2.4 for a precise definition) and norm $\|\cdot\|$ (with query access), the SVP is to find a shortest non-zero element of \mathcal{L} under $\|\cdot\|$. The CVP is the inhomogeneous version which takes in addition a target vector \mathbf{x} , and where the problem becomes to find an element of \mathcal{L} closest to \mathbf{x} under $\|\cdot\|$.

In Chapter 5, we study the SVP and CVP under arbitrary norms. A majority of the algorithms developed for SVP and CVP have focused on the setting of the ℓ_2 norm [82, 7, 118, 70, 64, 61, 91], and the techniques used therein have not generalized to other norms (without incurring large polynomial approximation factors). An important exception are the algorithms based on randomized sieving, a technique developed by Ajtai, Kumar and Sivakumar [2] to give the first randomized $2^{O(n)}$ time algorithm for the SVP under ℓ_2 . Subsequently, the AKS sieving approach was generalized to solve SVP and $(1 + \epsilon)$ -approximate CVP in any norm [3, 18, 5, 33], where the algorithms use $2^{O(n)}$ and $(1 + \frac{1}{\epsilon})^{O(n)}$ time, space and randomness respectively.

A natural question is whether there exists deterministic algorithms for general norm SVP and CVP with similar or better running times with respect to AKS. Our main contribution is to show that this is indeed possible. Formally, we give polynomial space Turing reductions from SVP and $(1 + \epsilon)$ -approximate CVP in any norm to CVP in the ℓ_2 -norm which perform $2^{O(n)}$ and $2^{O(n)}(1 + \frac{1}{\epsilon})^n$ calls to the CVP oracle and arithmetic operations respectively. Instantiating our reductions with the current fastest deterministic algorithm for CVP in ℓ_2 , i.e. the $O(2^{2n})$ time and $O(2^n)$ space algorithm of Micciancio and Voulgaris [91], we get $2^{O(n)}$ and $2^{O(n)}(1 + \frac{1}{\epsilon})^n$ time algorithms for SVP and $(1 + \epsilon)$ -CVP respectively under any norm which run in $O(2^n)$ space.

The main tool behind our reductions is a novel method for enumerating the lattice points inside any convex body, which relies on the construction of M-Ellipsoid coverings from the previous chapter. The second major ingredient, used within our $(1 + \epsilon)$ -CVP algorithm, is a new technique for “sparsifying” a lattice with respect to an arbitrary norm which approximately maintains the lattice’s metric structure.

The results in this chapter are based on joint work with Chris Peikert and Santosh Vempala [36], as well as subsequent extensions.

1.1.4 Integer Programming

The Integer Programming Problem (IP) is one of the foundational problems in operations research and computer science. IP is a highly effective modeling paradigm for discrete optimization problems that was introduced in the 1950’s. The first appearance of IP in the literature is the seminal paper of Dantzig, Fulkerson and Johnson [37] who used cutting plane methods to solve (by hand) a Traveling Salesman Problem (TSP) on 49 cities. The study of fully general IPs began with the work of Gomory [55] in 1958, who gave the first finite cutting plane algorithm for general problem. After the introduction of NP-Completeness, the binary version of IP appeared on Karp’s

original list of 21 NP-Complete problems. It remains one of the most well studied optimization problems even today.

Here we study the general integer version of the IP problem. The general form in which we consider the problem is as follows: given a convex set $K \subseteq \mathbb{R}^n$ by a separation oracle, decide whether $K \cap \mathbb{Z}^n \neq \emptyset$. We also study a natural optimization variant, that is, given a convex function $f : K \rightarrow \mathbb{R}$ (equipped with a subgradient oracle), either decide that $K \cap \mathbb{Z}^n = \emptyset$ or compute $\mathbf{y} \in K \cap \mathbb{Z}^n$ minimizing f .

Structure. In Chapter 6, we study the structure of the IP problem. The focus of our analysis is Kinchine’s flatness theorem in the geometry of numbers. The flatness theorem states that every convex body not containing integer points has integer width bounded by a function of dimension. More explicitly, if $K \subseteq \mathbb{R}^n$ is integer free, there exists a partition of \mathbb{Z}^n into parallel hyperplanes such that at most $f(n)$ of these hyperplanes intersect K , where $f(n)$ is function of dimension alone. The flatness theorem is one of the main structural results used in algorithms for solving general integer IP.

Due to its applications to IP, and its importance within the geometry of numbers, many proofs of the flatness theorem have been presented [76, 7, 80, 67, 10, 12, 13], each giving different asymptotic estimates on the function $f(n)$. The best current bounds are $f(n) = \Omega(n)$, and $f(n) = O(n^{\frac{4}{3}} \text{polylog}(n))$ due to Rudelon [113] using a reduction by Banaszczyk [12]. As our main contribution, we improve on a proof of Banaszczyk, which achieves a reduction with small explicit constants from bounding $f(n)$ to bounding the classical ℓ^* estimate in convex geometry. We remark that the best known asymptotic bounds on $f(n)$ are indeed derived in this fashion. Our main improvement is to avoid Banaszczyk’s reliance on Talagrand’s majorizing measure theorem (which results in huge unknown constants), by using a new geometric characterization of the smoothing parameter of a lattice.

The work in this Chapter is based on joint work with Kai-Min Chung, Feng Hao Liu, and Chris Peikert [27].

Algorithms. In Chapter 7, we study algorithms for the IP problem. The major algorithmic breakthroughs for general integer IP are due to Lenstra [84] and later Kannan [70], who respectively gave $2^{O(n^3)}$ and $2^{O(n)}n^{2.5n}$ time algorithms (focusing on the dependence on dimension) for IP when K is given by a system of linear inequalities (i.e. a polyhedron). These algorithms both relied on then recent advances in convex optimization, as well as novel insights in geometry of numbers (in particular, the flatness theorem).

Our main contributions are faster algorithms for both feasibility and optimization variants of IP. The first algorithm, based on a framework of Lenstra, solves IP feasibility problems in $2^{O(n)}(n^{\frac{4}{3}} \text{polylog}(n))^n$ time and $O(2^n)$ space. The second algorithm, based on a framework of Kannan, solves IP feasibility in $2^{O(n)}n^n$ time and $O(2^n)$ space. Our last algorithm, directly solves integer optimization problems. In expected $2^{O(n)}n^n$ time and $O(2^n)$ space algorithm it can minimize any convex function over the integer points in a bounded convex set. Here, our IP algorithms crucially rely on both the lattice algorithms and lattice point enumeration techniques developed in Chapter 5.

The results in this Chapter are based on joint work with Chris Peikert and Santosh Vempala [36], as well as subsequent extensions.

CHAPTER II

MATHEMATICAL BACKGROUND

In this chapter, we present the basic concepts and notational conventions used throughout this thesis. It is intended as a reference only, and may be skipped upon first reading. Throughout the thesis, we shall refer the reader to relevant sections of the background chapter to elucidate the used notation and concepts.

2.1 *Basics*

We denote \mathbb{C} the set of complex numbers, \mathbb{R} the set of real numbers, \mathbb{Q} the set of rational numbers, and \mathbb{Z} the set of integers. We write $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ for the extended reals (∞ denotes infinity). The sets $\bar{\mathbb{R}}_+, \mathbb{R}_+, \mathbb{Q}_+, \mathbb{Z}_+$ denote the corresponding non-negative versions. We define the natural numbers \mathbb{N} , as the set of positive integers. For $n \in \mathbb{N}$, we denote $[n] = \{1, \dots, n\}$. For sets S, T we denote their cross product $S \times T = \{(s, t) : s \in S, t \in T\}$. For $n \geq 1$, we define $S^n = \overbrace{S \times S \times \dots \times S}^{n \text{ times}}$.

For some notational conventions, we write vectors in bold and scalars in regular font, i.e. $\mathbf{x} \in \mathbb{R}^n$ and $x \in \mathbb{R}$. As a further convention, we let $\mathbf{0}$ denote either the all zero vector or matrix for the ambient space (which will be clear from context). We write \emptyset to denote the emptyset.

Functions and Sets. Let $f : X \rightarrow Y$ be a function from X to Y . For $S \subseteq X$ we denote the image of S under f as $f(S) = \{f(\mathbf{s}), \mathbf{s} \in S\}$, and for $T \subseteq Y$ we denote the inverse image of T under f as $f^{-1}(T) = \{\mathbf{x} \in X : f(\mathbf{x}) \in T\}$. For functions $f : X \rightarrow Y$, and $g : Y \rightarrow Z$, we define the composition $g \circ f : X \rightarrow Z$ by the relation $(g \circ f)(\mathbf{x}) = g(f(\mathbf{x}))$. We say that $f : X \rightarrow Y$ is injective if for $\mathbf{x}, \mathbf{z} \in X$, $\mathbf{x} \neq \mathbf{z}$, $f(\mathbf{x}) \neq f(\mathbf{z})$, and we say that f is surjective if $\forall \mathbf{y} \in Y$ there exists $\mathbf{x} \in X$ such that

$f(\mathbf{x}) = \mathbf{y}$ (i.e. $f(X) = Y$). Lastly, f is bijective if it is both surjective and injective.

We say that a set X is countable if there exists a mapping $f : \mathbb{N} \rightarrow X$ such that f is surjective.

For sets $A, B \subseteq \mathbb{R}^n$ we define the Minkowski sum of A and B as

$$A + B = \{\mathbf{x} + \mathbf{y} : \mathbf{x} \in A, \mathbf{y} \in B\}.$$

For a vector $\mathbf{t} \in \mathbb{R}^n$, we define $\mathbf{t} + A = \{\mathbf{t}\} + A$ for notational convenience. For a scalar $s \in \mathbb{R}$, we define $sA = \{s\mathbf{a} : \mathbf{a} \in A\}$. For a set of scalars $S \subseteq \mathbb{R}$ and vector $\mathbf{x} \in \mathbb{R}^n$, we define $S\mathbf{x} = \{s\mathbf{x} : s \in S\}$. The above definitions extend in the natural way for any spaces supporting the requisite addition and multiplication operations.

For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we denote the line (or interval) between \mathbf{x} and \mathbf{y} as

$$[\mathbf{x}, \mathbf{y}] = \{\alpha\mathbf{x} + (1 - \alpha)\mathbf{y} : 0 \leq \alpha \leq 1\}.$$

We denote $(\mathbf{x}, \mathbf{y}) = [\mathbf{x}, \mathbf{y}] \setminus \{\mathbf{x}\}$ ((\mathbf{y}, \mathbf{x}) is defined analogously), and $(\mathbf{x}, \mathbf{y}) = [\mathbf{x}, \mathbf{y}] \setminus \{\mathbf{x}, \mathbf{y}\}$, the half-open, and open interval between \mathbf{x} and \mathbf{y} respectively. We note that for numbers $x, y \in \mathbb{R}$, $x \leq y$, from the above definition, the set $[x, y] = \{z \in \mathbb{R} : x \leq z \leq y\}$ ($[x, y)$ and (x, y) are similarly understood).

For a real number $x \in \mathbb{R}$, we define $\lfloor x \rfloor$ as the greatest integer less than or equal to x , and $\lceil x \rceil$ as the least integer greater than or equal to x . Note that $x - \lfloor x \rfloor \in [0, 1)$. For a set $S \subseteq \bar{\mathbb{R}}$, we denote $\sup S$ the smallest element $x \in \bar{\mathbb{R}}$ satisfying $x \geq y \quad \forall y \in S$, and $\inf S$ the largest element $x \in \bar{\mathbb{R}}$ satisfying $x \leq y \quad \forall y \in S$. If $(\inf) \sup S \in S$, i.e. S contains its (infimum) supremum, we write $(\inf) \sup S = (\min) \max S$. For a function $f : S \rightarrow \mathbb{R}$, we define $\arg \max_{x \in S} f(x) = \{x \in S : f(x) = \max\{f(x) : x \in S\}\}$ and define $\arg \min_{x \in S} f(x)$ analogously.

Modular Arithmetic. For integers $a, b \in \mathbb{Z}$, $m \in \mathbb{N}$, we define the equivalence relation $a \equiv b \pmod{m} \Leftrightarrow m \mid a - b$ (m divides $a - b$). We write the set of equivalence classes as $\mathbb{Z}_m = \{a + m\mathbb{Z} : a \in \mathbb{Z}\}$ (i.e. the integers mod m), where $|\mathbb{Z}_m| = m$. Here

one can check that for $a, b, c, d \in \mathbb{Z}$ such that $a \equiv b \pmod{m}$ and $c \equiv d \pmod{m}$ then $a + c \equiv b + d \pmod{m}$ and $ac \equiv bd \pmod{m}$, and hence both addition and multiplication are well defined on \mathbb{Z}_m . Therefore \mathbb{Z}_m is a ring. If $p \in \mathbb{N}$ is prime (has only 1 and p as divisors), then for every $a \not\equiv 0 \pmod{p}$, there exists $b \in \mathbb{Z}$ such that $ab \equiv 1 \pmod{p}$. In this case all non-zero elements of \mathbb{Z}_p have multiplicative inverses, and hence \mathbb{Z}_p is a field.

For a treatment of groups, fields and modular arithmetic, the reader may consult [42].

Graphs. A graph G is a tuple (V, E) , denoting a vertex set V , which is either countable or finite (often identified as a subset of \mathbb{N}), and edge set $E \subseteq \{\{x, y\} : x \in V, y \in V\}$, i.e. a subset of pairs of vertices. A directed graph G is tuple (V, E) , where V again denotes vertices, and where $E \subseteq V^2$, i.e. are ordered pairs of vertices (so (v, w) and (w, v) are distinct for distinct $v, w \in V$). A graph $H = (V_1, E_1)$ is a subgraph of a graph $G = (V, E)$ if $V_1 \subseteq V$ and $E_1 \subseteq E$. H is an induced subgraph on G if $E_1 = E \cap \{\{v, w\} : v, w \in V_1\}$ (i.e. all edges on the vertices of V_1 in G are present). We analogously define the subgraph and induced subgraph relations for directed graphs. Lastly, for a directed graph $H = (V_1, E_1)$ and the (undirected) graph $G = (V, E)$, we say that H is a directed subgraph of G if $V_1 \subseteq V$ and $E'_1 = \{\{v, w\} : (v, w) \in E_1 \text{ or } (w, v) \in E_1\} \subseteq E$ (i.e. the undirected “version” of H is a subgraph of G).

For a (directed) graph $G = (V, E)$, the vertices v and w are adjacent, which we write $v \sim w$, if $((v, w)) \{v, w\} \in E$. An ordered list of distinct vertices $P = (v_1, \dots, v_k)$ in a (directed) graph G is a (directed) path if $v_i \sim v_{i+1}$ for $i \in [k - 1]$. An ordered list $C = (v_1, \dots, v_k, v_1)$, where each v_i are distinct and $k \geq 2$, is a cycle if $v_i \sim v_{i+1}$ for $i \in [k - 1]$ and $v_k \sim v_1$. A graph G is connected if for all $v, w \in V$, there is a path starting at v and ending at w . G is disconnected if G is not connected.

A connected component of a graph G , is a connected induced subgraph of G that is vertex maximal (i.e. adding any additional vertex would make it disconnected).

For a graph $G = (V, E)$, and vertex $v \in V$, we denote the edges incident to v as $E(v) = \{\{v, w\} : w \in V, \{v, w\} \in E\}$. For a directed graph $G = (V, E)$, and vertex $v \in V$, we denote the in-edges at v as $E^-(v) = \{(w, v) : w \in V, (w, v) \in E\}$ and the out-edges at v as $E^+(v) = \{(v, w) : w \in V, (v, w) \in E\}$. We say that $v \in V$ is a sink, if $|E^+(v)| = 0$.

A graph G is a forest if G is acyclic (i.e. G contains no cycles). G is a tree if G is both acyclic and connected. A directed graph G is a directed rooted forest if the underlying undirected graph (i.e. forced all edges to be undirected) is a forest and each connected component contains a unique sink vertex. Under this definition, note that every vertex $v \in V$ must have $|E^+(v)| \leq 1$. Furthermore, from every $v \in V$ there exists a unique path along out-edges leading to a sink vertex of G . Lastly, G is a directed rooted tree if the underlying undirected graph is a tree and G has a unique sink vertex.

For a treatment of graphs and their properties, the reader may consult [39].

2.1.1 Linear Spaces

For a thorough treatment of linear spaces and their properties, the reader may consult [121, 54].

$W \subseteq \mathbb{R}^n$ is a linear subspace if for $\mathbf{x}, \mathbf{y} \in W$ and $a, b \in \mathbb{R}$, $a\mathbf{x} + b\mathbf{y} \in W$. An affine subspace $T \subseteq \mathbb{R}^n$ is set of the form $W + \mathbf{x}$, where $W \subseteq \mathbb{R}^n$ is a linear subspace and $\mathbf{x} \in \mathbb{R}^n$.

Vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n$ are linearly independent if $\sum_{i=1}^k a_i \mathbf{b}_i = \mathbf{0} \Leftrightarrow a_1 = \dots = a_k = 0$. $\mathbf{b}_1, \dots, \mathbf{b}_k$ form a basis for a linear subspace $W \subseteq \mathbb{R}^n$, if they are linear independent and $\mathbb{R}\mathbf{b}_1 + \dots + \mathbb{R}\mathbf{b}_k = W$. From basic linear algebra, we have that every linear subspace W admits a basis, and every basis of W is composed of the

same number of vectors. We define $\dim(W)$, the dimension of the subspace W , as the number of vectors in any basis of W . For an affine space $T = W + \mathbf{x}$, we define $\dim(T) = \dim(W)$. An affine space $T \subseteq \mathbb{R}^n$ is a hyperplane if $\dim(T) = n - 1$.

Let $A \subseteq \mathbb{R}^n$. We denote $\text{span}(A)$ the linear span of A , i.e. the smallest linear subspace containing A , and $\text{aff}(A)$ the affine hull of A , i.e. the smallest affine subspace containing A . We define $\dim(A) = \dim(\text{aff}(A))$.

Matrices. We let $\mathbb{R}^{n \times m}$, integers $n, m \geq 1$, denote the set of matrices with n rows and m columns with entries in \mathbb{R} . For $A \in \mathbb{R}^{n \times m}$, we write $(A)_{ij}$ for the ij^{th} entry of A . In this notation we identify \mathbb{R}^n above with $\mathbb{R}^{n \times 1}$, i.e. column vectors with n rows. Furthermore for $\mathbf{x} \in \mathbb{R}^n$, we write $(\mathbf{x})_i$ (or \mathbf{x}_i when there is no confusion with other subscripts) for the i^{th} entry of \mathbf{x} . We write $A = (\mathbf{a}_1, \dots, \mathbf{a}_m) \in \mathbb{R}^{n \times m}$, to denote the $n \times m$ matrix whose i^{th} column is $\mathbf{a}_i \in \mathbb{R}^n$, $1 \leq i \leq m$.

We denote $\mathbf{e}_1, \dots, \mathbf{e}_n \in \mathbb{R}^n$, the standard unit vectors, which satisfy $(\mathbf{e}_i)_j = 1$ if $i = j$ and 0 otherwise, for $i, j \in [n]$.

We define $A^T \in \mathbb{R}^{m \times n}$, the transpose of A , by relation $(A^T)_{ij} = (A)_{ji}$. For $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{m \times p}$, we define the matrix product $AB \in \mathbb{R}^{n \times p}$ by the relation

$$(AB)_{ij} = \sum_{l=1}^m (A)_{il}(B)_{lj} \quad \text{for } i \in [n], j \in [p].$$

For $S \subseteq \mathbb{R}^m$, we define the image of S under A as $A(S) = \{A\mathbf{x} : \mathbf{x} \in S\}$ (we also write AS when there is no confusion). We define the kernel of the matrix A by $\text{kern}(A) = \{\mathbf{x} \in \mathbb{R}^m : A\mathbf{x} = \mathbf{0}\}$.

Define $I_n \in \mathbb{R}^{n \times n}$, the $n \times n$ identity matrix by the rule $(I_n)_{ij} = 1$ if $i = j$ and 0 otherwise. For $A \in \mathbb{R}^{n \times n}$, we define A^{-1} the inverse of A (when it exists) to be the matrix satisfying $A^{-1}A = AA^{-1} = I_n$. We say that A is invertible (or non-singular) if A^{-1} exists. A matrix $A \in \mathbb{R}^{n \times n}$ is orthogonal if $A^T = A^{-1}$, i.e. $A^T A = I_n$. We define the trace of A by $\text{tr}(A) = \sum_{i=1}^n (A)_{ii}$. For the trace operator, we have the relation $\text{tr}(AB) = \text{tr}(BA)$ for matrices $A, B \in \mathbb{R}^{n \times n}$.

A map $T : V \rightarrow W$, where $V \subseteq \mathbb{R}^n$ and $W \subseteq \mathbb{R}^m$ are linear subspaces, is a linear transformation (or linear map) if $T(a\mathbf{x} + b\mathbf{y}) = aT(\mathbf{x}) + bT(\mathbf{y})$ for $a, b \in \mathbb{R}$ and $\mathbf{x}, \mathbf{y} \in W$. Given bases $\mathbf{v}_1, \dots, \mathbf{v}_k$ and $\mathbf{w}_1, \dots, \mathbf{w}_l$ for V and W respectively, we can associate a matrix to $A \in \mathbb{R}^{l \times k}$ to the transformation T by the relation

$$T\mathbf{v}_i = \sum_{j=1}^l (A)_{ji} \mathbf{w}_j \text{ for } i \in [k].$$

Given $\mathbf{v} \in V$, letting $\mathbf{a} \in \mathbb{R}^k$ denote the unique vector satisfying $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_k)\mathbf{a}$, we recover $T(\mathbf{v})$ as follows:

$$T(\mathbf{v}) = (\mathbf{w}_1, \dots, \mathbf{w}_l)A\mathbf{a}.$$

Hence, after choosing a basis, linear transformations can be expressed via appropriate matrix multiplications.

Norms. A function $\mathbf{p} : \mathbb{R}^n \rightarrow \mathbb{R}_+$ defines a norm on \mathbb{R}^n if $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, and $t \in \mathbb{R}_+$, we have that

(1) $\mathbf{p}(\mathbf{x} + \mathbf{y}) \leq \mathbf{p}(\mathbf{x}) + \mathbf{p}(\mathbf{y})$ (triangle inequality)

(2) $\mathbf{p}(t\mathbf{x}) = t\mathbf{p}(\mathbf{x})$ (positive homogeneity)

(3) $\mathbf{p}(\mathbf{x}) = \mathbf{p}(-\mathbf{x})$ (symmetry)

(4) $\mathbf{p}(\mathbf{x}) = 0 \Leftrightarrow \mathbf{x} = \mathbf{0}$ (non-degeneracy)

If \mathbf{p} satisfies all the above condition except symmetry, we call \mathbf{p} an asymmetric norm. When we refer to general norms, we include both symmetric and asymmetric norms. We define the unit ball of the norm \mathbf{p} to be $B_{\mathbf{p}} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{p}(\mathbf{x}) \leq 1\}$.

Define the ℓ_p norms on \mathbb{R}^n , $p \geq 1$, as $\|\mathbf{x}\|_p = (\sum_{i=1}^n |\mathbf{x}_i|^p)^{\frac{1}{p}}$. That $\|\cdot\|_p$, $p \geq 1$, defines a norm on \mathbb{R}^n is classical fact from analysis. For convenience we generally write $\|\mathbf{x}\|$ for $\|\mathbf{x}\|_2$, i.e. the standard euclidean norm. We note that $\|\mathbf{x}\|_{\infty} = \max_{i \in [n]} |\mathbf{x}_i|$.

Let $B_p^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_p \leq 1\}$, $p \geq 1$, denote the ℓ_p ball in \mathbb{R}^n . Here B_p^n is the unit ball of the norm $\|\cdot\|_p$. We note that $B_\infty^n = [-1, 1]^n$ and B_2^n is the unit euclidean ball in \mathbb{R}^n .

For a matrix $A \in \mathbb{R}^{n \times m}$, we denote its operator norm as $\|A\|_2 = \max_{\mathbf{x} \in B_2^n} \|A\mathbf{x}\|_2$, and its Frobenius norms as $\|A\|_F = \sqrt{\sum_{ij} (A)_{ij}^2}$.

Determinant. We define determinant $\det : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ as the unique function on $n \times n$ matrices satisfying:

(1) Multilinearity: For $\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{x} \in \mathbb{R}^n$, $s, t \in \mathbb{R}$ and $1 \leq i \leq n$,

$$\det(\mathbf{a}_1, \dots, s\mathbf{a}_i + t\mathbf{x}, \dots, \mathbf{a}_n) = s \det(\mathbf{a}_1, \dots, \mathbf{a}_n) + t \det(\mathbf{a}_1, \dots, \mathbf{a}_{i-1}, \mathbf{x}, \mathbf{a}_{i+1}, \dots, \mathbf{a}_n).$$

(2) Anti-symmetry: For $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^n$, $\mathbf{a}_i = \mathbf{a}_j$ for any $i \neq j$ implies that

$$\det(\mathbf{a}_1, \dots, \mathbf{a}_n) = 0.$$

(3) Normalization: $\det(I_n) = 1$.

From elementary linear algebra, for $A \in \mathbb{R}^{n \times n}$ we have that $\det(A) \neq 0$ iff A is non-singular. Furthermore, for $A, B \in \mathbb{R}^{n \times n}$, we have that $\det(AB) = \det(A)\det(B)$ and $\det(A) = \det(A^T)$.

Positive Semi-Definite Matrices. A matrix $A \in \mathbb{R}^{n \times n}$ is symmetric if $A^T = A$. A is positive (semi-)definite, which we write $A \succ (\succeq) \mathbf{0}$, if A is symmetric and $\forall \mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$, $\mathbf{x}^T A \mathbf{x} > (\geq) 0$. We note that relation \succeq (\succ) yields a partial ordering on the space of symmetric matrices, i.e. for $A, B \in \mathbb{R}^{n \times n}$ symmetric, $A \succeq (\succ) B \Leftrightarrow (A - B) \succeq (\succ) \mathbf{0}$. Furthermore, the ordering is stable under the following operation, if $B \in \mathbb{R}^{n \times n}$ and $A \succeq \mathbf{0}$ then $B^T A B \succeq \mathbf{0}$. Furthermore if B is non-singular then $A \succ (\succeq) \mathbf{0} \Leftrightarrow B^T A B \succ (\succeq) \mathbf{0}$. From here, we see that $A \succ \mathbf{0}$ if and only if $A^{-1} \succ \mathbf{0}$.

Inner Products. An inner product $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies:

- (1) For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$ (symmetry).
- (2) For $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^n$, $a, b \in \mathbb{R}$, $\langle a\mathbf{x} + b\mathbf{y}, \mathbf{z} \rangle = a \langle \mathbf{x}, \mathbf{z} \rangle + b \langle \mathbf{y}, \mathbf{z} \rangle$ (bilinearity)
- (3) For $\mathbf{x} \in \mathbb{R}^n$, $\langle \mathbf{x}, \mathbf{x} \rangle > 0$ if $\mathbf{x} \neq \mathbf{0}$ (non-degeneracy)

An inner product $\langle \cdot, \cdot \rangle$ satisfies the Cauchy-Schwarz inequality: $\langle \mathbf{x}, \mathbf{y} \rangle^2 \leq \langle \mathbf{x}, \mathbf{x} \rangle \langle \mathbf{y}, \mathbf{y} \rangle$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. Furthermore the inequality holds with equality iff \mathbf{x} and \mathbf{y} are collinear (i.e. $\mathbf{x} = t\mathbf{y}$ for some $t \in \mathbb{R}$).

An inner-product $\langle \cdot, \cdot \rangle$ induces a matrix $A \in \mathbb{R}^{n \times n}$ by the relation $(A)_{ij} = \langle \mathbf{e}_i, \mathbf{e}_j \rangle$. From here it is easy to verify that $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T A \mathbf{y}$. Furthermore, $\langle \cdot, \cdot \rangle$ is an inner product if and only if the induced matrix A is positive definite, i.e. $A \succ 0$. Henceforth, for a matrix $A \succ 0$, $A \in \mathbb{R}^{n \times n}$, we denote the inner product $\langle \mathbf{x}, \mathbf{y} \rangle_A = \mathbf{x}^T A \mathbf{y}$. When no matrix A is specified, we shall intend A to be the identity, i.e. the standard inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T \mathbf{y}$.

We define $\|\mathbf{x}\|_A = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_A} = \sqrt{\mathbf{x}^T A \mathbf{x}}$ to be the norm induced by the inner product $\langle \cdot, \cdot \rangle_A$. We refer to $\|\cdot\|_A$ as an ellipsoidal norm. That $\|\cdot\|_A$ is a norm follows directly from the Cauchy-Schwarz inequality. We note that the standard ℓ_2 norm is induced by the standard inner product $\langle \cdot, \cdot \rangle$.

Vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n$ are orthogonormal if $\langle \mathbf{b}_i, \mathbf{b}_j \rangle = 0$ if $i \neq j$, and are orthonormal if they additionally satisfy $\|\mathbf{b}_i\|_2 = 1 \quad \forall i \in [n]$.

For a linear subspace $W \subseteq \mathbb{R}^n$, we define its orthogonal complement $W^\perp = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \quad \forall \mathbf{y} \in W\}$. From basic linear algebra, we have that $W \cap W^\perp = \{\mathbf{0}\}$ and $W + W^\perp = \mathbb{R}^n$. For notional convenience, for a vector $\mathbf{x} \in \mathbb{R}^n$, we write \mathbf{x}^\perp for $\text{span}(\mathbf{x})^\perp$.

For a linear subspace $W \subseteq \mathbb{R}^n$, we define the orthogonal projection $\pi_W : \mathbb{R}^n \rightarrow W$ as the linear transformation defined by the relation $\pi_W(\mathbf{x}) = \bar{\mathbf{x}}$ where $\bar{\mathbf{x}} \in W$ is the unique vector satisfying $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \bar{\mathbf{x}}, \mathbf{y} \rangle \quad \forall \mathbf{y} \in W$. For an affine subspace $T = W + \mathbf{c}$, $\mathbf{c} \in \mathbb{R}^n$, we define the orthogonal projection $\pi_T : \mathbb{R}^n \rightarrow T$ to be the affine map

$\mathbf{x} \rightarrow \pi_W(\mathbf{x}) + \mathbf{t}$, where \mathbf{t} is the unique vector in $T \cap W^\perp$. The orthogonal projection map can also be rephrased as follows: for an affine subspace $T \subseteq \mathbb{R}^n$, the orthogonal projection π_T corresponds with the map $\mathbf{x} \rightarrow \arg \min_{\mathbf{y} \in T} \|\mathbf{x} - \mathbf{y}\|_2$.

Gram Schmidt Orthogonalization. For vectors linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n$, we define their gram schmidt orthogonalization (or gram schmidt vectors) $\mathbf{b}_1^*, \dots, \mathbf{b}_k^*$ as follows: $\mathbf{b}_1^* = \mathbf{b}_1$ and $\mathbf{b}_i^* = \mathbf{b}_i - \sum_{j=1}^{i-1} \frac{\langle \mathbf{b}_i, \mathbf{b}_j^* \rangle}{\langle \mathbf{b}_j^*, \mathbf{b}_j^* \rangle} \mathbf{b}_j^*$. We note that the vectors $\mathbf{b}_1^*, \dots, \mathbf{b}_k^*$ are orthogonal. Given vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$ we define the map $\pi_i : \mathbb{R}^n \rightarrow \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{i-1})^\perp$, $1 \leq i \leq k$, as the orthogonal projection map onto $\text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{i-1})^\perp$. Under this definition, the gram schmidt vectors satisfy $\mathbf{b}_i^* = \pi_i(\mathbf{b}_i) \quad \forall i \in [k]$.

For $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n$, we define the gram matrix $\text{Gram}(\mathbf{b}_1, \dots, \mathbf{b}_k) = G \in \mathbb{R}^{k \times k}$ by the relation $(G)_{ij} = \langle \mathbf{b}_i, \mathbf{b}_j \rangle$. Equivalently, letting $B = (\mathbf{b}_1, \dots, \mathbf{b}_k)$, we have that $G = B^T B$. It is easy to see that $G \succeq \mathbf{0}$, and that $G \succ \mathbf{0}$ if and only if $\mathbf{b}_1, \dots, \mathbf{b}_k$ are linearly independent. Furthermore, we have that $\det(G) = \prod_{i=1}^k \|\mathbf{b}_i^*\|_2^2$, where $\mathbf{b}_1^*, \dots, \mathbf{b}_k^*$ are the gram schmidt vectors of $\mathbf{b}_1, \dots, \mathbf{b}_k$.

Eigen Values. For a matrix $A \in \mathbb{C}^{n \times n}$, its eigen values are the roots of the characteristic polynomial $\det(A - \lambda I_n)$ (univariate polynomial in the variable λ). In particular, A has an eigen value $\lambda \in \mathbb{C}$ if and only if there is an associated eigen vector $\mathbf{v} \in \mathbb{C}^n$ satisfying $A\mathbf{v} = \lambda\mathbf{v}$. We note that $\det(A - \lambda I_n)$ is a polynomial of degree n with complex coefficients, and hence by the fundamental theorem of Algebra, has exactly n complex roots $\lambda_1, \dots, \lambda_n$ (counting multiplicities). Furthermore, we have that $\det(A) = \prod_{i=1}^n \lambda_i$ and $\text{tr}(A) = \sum_{i=1}^n \lambda_i$.

If $A \in \mathbb{R}^{n \times n}$ is symmetric ($A^T = A$), then A admits a basis of real orthonormal eigen vectors $\mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^n$ with associated real eigen values $\lambda_1, \dots, \lambda_n \in \mathbb{R}$. Given such a basis, we can write $A = \sum_{i=1}^n \lambda_i \mathbf{b}_i \mathbf{b}_i^T$. Furthermore, $A \succeq (\succ) \mathbf{0}$ if and only if $\lambda_1, \dots, \lambda_n \geq (>) 0$. Also, if $A \succeq \mathbf{0}$ there is a unique matrix $A^{\frac{1}{2}}$ (the square root of A),

satisfying $A^{\frac{1}{2}} \succcurlyeq \mathbf{0}$ and $(A^{\frac{1}{2}})^2 = A$. In particular, $A^{\frac{1}{2}} = \sum_{i=1}^n \lambda_i^{\frac{1}{2}} \mathbf{b}_i \mathbf{b}_i^T$.

2.1.2 Elementary Analysis

For a thorough treatment of the point set topology, continuity and differentiation, the reader may consult [115, 22].

Point Set Topology. A set $S \subseteq \mathbb{R}^n$ is open, if for all $\mathbf{x} \in S$ there exists $\epsilon > 0$ such that $\mathbf{x} + \epsilon B_2^n \subseteq S$. $S \subseteq \mathbb{R}^n$ is closed if and only if $\mathbb{R}^n \setminus S$ is open. S is bounded if there exists $R < \infty$ such that $S \subseteq RB_2^n$. The closure $\text{cl}(S)$ of S is the smallest closed set containing S , and the interior $\text{int}(S)$ of S is the largest open set contained in S . We denote the boundary of S as $\text{bd}(S) = \text{cl}(S) \setminus \text{int}(S)$ (we also occasionally write ∂S for $\text{bd}(S)$). We say that S is full dimensional if $\text{int}(S) \neq \emptyset$.

For an affine subspace $T \subseteq \mathbb{R}^n$, we say that $S \subseteq T$ is open with respect to T if $\forall \mathbf{x} \in S$ there exists $\epsilon > 0$ such that $(\mathbf{x} + \epsilon B_2^n) \cap T \subseteq S$. Similarly, $S \subseteq T$ is closed in T if $T \setminus S$ is open in T . We define the interior of S with respect to T , as the largest open set in T contained in S , and the closure of S with respect to T , as the smallest closed set in T containing S . The boundary of S with respect to T is defined analogously.

For $S \subseteq \mathbb{R}^n$, we define the relative interior $\text{relint}(S)$, as the interior of S with respect to affine subspace $\text{aff}(S)$. Similarly, we define the relative boundary $\text{relbd}(S)$, as the boundary of S with respect to the affine subspace $\text{aff}(S)$.

For a sequence $(\mathbf{x}_i)_{i=1}^{\infty} \subseteq \mathbb{R}^n$, we write that $\lim_{i \rightarrow \infty} \mathbf{x}_i = \mathbf{x} \in \mathbb{R}^n$ (i.e. \mathbf{x} is the limit of the sequence $\mathbf{x}_1, \mathbf{x}_2, \dots$), if for all $\epsilon > 0$, there exists $N_0 \in \mathbb{N}$, such that for all $i \geq N_0$, $\|\mathbf{x} - \mathbf{x}_i\| \leq \epsilon$. We may also define limits for sequences of functions. Let f_1, f_2, \dots denote a countable sequence of functions, where $f_i : U \rightarrow V$ for $i \geq 1$, where $U \subseteq \mathbb{R}^n$ and $V \subseteq \mathbb{R}^m$. For $f : U \rightarrow V$, we write $\lim_{i \rightarrow \infty} f_i = f$, if $\forall \mathbf{x} \in U$ we have that $\lim_{i \rightarrow \infty} f_i(\mathbf{x}) = f(\mathbf{x})$.

A set $C \subseteq \mathbb{R}^n$ is compact if for any collection $\mathcal{U} = \{U_i : i \geq 1\}$ of open sets in \mathbb{R}^n such that $C \subseteq \cup_{U \in \mathcal{U}} U$ there exists a finite subcollection $\mathcal{U}' \subseteq \mathcal{U}$, $|\mathcal{U}'| < \infty$ such that

$C \subseteq \cup_{U \in \mathcal{U}}$. We now state some fundamental equivalences about compact sets in \mathbb{R}^n which will be useful in the segway.

Theorem 2.1.1. *The following statements are equivalent:*

- (1) $C \subseteq \mathbb{R}^n$ is compact.
- (2) $C \subseteq \mathbb{R}^n$ is closed and bounded.
- (3) For any sequence $(\mathbf{x}_i)_{i=1}^{\infty} \subseteq C$, there exists a subsequence $(\mathbf{x}_{s_i})_{i=1}^{\infty}$ such that the limit $\lim_{i \rightarrow \infty} \mathbf{x}_{s_i} = \bar{\mathbf{x}}$ exists and $\bar{\mathbf{x}} \in C$.

Continuity and Differentiation. Let $f : U \rightarrow V$ be a function from $U \subseteq \mathbb{R}^n$ to $V \subseteq \mathbb{R}^m$. f is continuous at a point $\mathbf{x} \in U$, if $\forall \epsilon > 0$ there exists $\delta > 0$ such that $\forall \mathbf{y} \in U$, $\|\mathbf{y} - \mathbf{x}\| \leq \delta$, we have that $\|f(\mathbf{x}) - f(\mathbf{y})\| \leq \epsilon$. f is continuous if for all $\mathbf{x} \in U$, f is continuous at \mathbf{x} .

Let $f : U \rightarrow V$ be function from $U \subseteq \mathbb{R}^n$ to $V \subseteq \mathbb{R}^m$, where U, V are open sets. f is differentiable at a point $\mathbf{x} \in U$, if there exists a matrix $T_{\mathbf{x}} \in \mathbb{R}^{n \times m}$ such that $\forall \epsilon > 0$ there exist $\delta > 0$ such that $\forall \mathbf{h} \in \mathbb{R}^n$, $\|\mathbf{h}\|_2 \leq \delta$, we have that

$$\frac{\|f(\mathbf{x} + \mathbf{h}) - f(\mathbf{x}) - T_{\mathbf{x}}\mathbf{h}\|}{\|\mathbf{h}\|} \leq \epsilon$$

f is differentiable (on $S \subseteq U$) if $\forall \mathbf{x} \in U$ ($\forall \mathbf{x} \in S$), f is differentiable at \mathbf{x} . Here it is easy to verify that if f is differentiable, then it is also continuous. If $U, V \subseteq \mathbb{R}$ (i.e. f is 1 dimensional), then for $x \in U$, we write $f'(x) \in \mathbb{R}$ for the differential T_x of f at x .

A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is monotone non-decreasing (non-increasing) if for $x, y \in \mathbb{R}$, $x \leq y$, $f(x) \leq (\geq) f(y)$. If f is differentiable, then f is monotone non-decreasing (non-increasing) if and only if $\forall x \in \mathbb{R}$, $f'(x) \geq (\leq) 0$.

2.1.3 Probability and Measure

For a detailed expositions on measure theory, integration, and probability the reader may consult the following references [114, 16, 60].

For a domain Ω , we let 2^Ω (the power set of Ω) denote the set of all subsets of Ω . $\mathcal{A} \subseteq 2^\Omega$ is a σ -algebra on Ω , if \mathcal{A} satisfies (1) $\mathcal{A} \neq \emptyset$, (2) $A \in \mathcal{A} \Rightarrow \Omega \setminus A \in \mathcal{A}$ and (3) $A_1, A_2, \dots \in \mathcal{A} \Rightarrow \bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$. A measure $\mu : \mathcal{A} \rightarrow \overline{\mathbb{R}}_+$ is function satisfying

$$(1) \quad \mu(\emptyset) = 0.$$

(2) For A_1, A_2, \dots a countable collection of pairwise disjoint sets:

$$\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i)$$

We denote the triple $(\Omega, \mathcal{A}, \mu)$ a measure space. If $\mu(\Omega) = 1$, we say that $(\Omega, \mathcal{A}, \mu)$ is a probability measure. A measure space $(\Omega, \mathcal{A}, \mu)$ is σ -finite if Ω can be expressed as the union of countably many sets in \mathcal{A} of finite measure.

Let $(\Omega_1, \mathcal{A}_1, \mu_1), (\Omega_2, \mathcal{A}_2, \mu_2)$ denote two σ -finite measure spaces. We define the product measure $\mu_1 \times \mu_2$ on $\Omega_1 \times \Omega_2$, as the unique measure defined on the σ -algebra induced by $\mathcal{A}_1 \times \mathcal{A}_2$ satisfying $(\mu_1 \times \mu_2)(A_1 \times A_2) = \mu_1(A_1) \times \mu_2(A_2)$.

We define the Borel sets in \mathbb{R}^n , which we write \mathcal{B}^n , as the smallest σ -algebra containing the open sets in \mathbb{R}^n . Let $(\Omega, \mathcal{A}, \mu)$ denote a measure space. A function $f : \Omega \rightarrow \mathbb{R}^n$ is measurable, if $f^{-1}(A) \in \mathcal{A} \quad \forall A \in \mathcal{B}^n$ (i.e. the Borel sets). A function $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is measurable if $g^{-1}(A) \in \mathcal{B}^n$ for $A \in \mathcal{B}^m$. Here we have that the composition $g \circ f : \Omega \rightarrow \mathbb{R}^m$ is also measurable under our definitions. If $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuous, it is again easy to check that g is measurable. For measurable functions $f, g : \Omega \rightarrow \mathbb{R}^n$, the functions $f+g, f \cdot g$, and (f, g) (i.e. $(f, g)(\mathbf{x}) = (f(\mathbf{x}), g(\mathbf{x}))$) are also measurable. Lastly, for a convergent sequence of measurable functions $f_i : \Omega \rightarrow \mathbb{R}^n$, $i \geq 1$, the function $f = \lim_{i \rightarrow \infty} f_i$ is measurable.

Let vol_n (n dimensional volume), denote the Lebesgue measure on \mathbb{R}^n . Restricted to \mathcal{B}^n , the Lebesgue measure vol_n is unique translation invariant measure on \mathbb{R}^n (i.e. $\text{vol}_n(A + \mathbf{x}) = \text{vol}_n(A)$ for $A \in \mathcal{B}^n, \mathbf{x} \in \mathbb{R}^n$) satisfying $\text{vol}_n([0, 1]^n) = 1$. We define a subset $S \subseteq \mathbb{R}^n$ to be measurable if $S \in \mathcal{B}^n$ (i.e. S is Borel-measurable¹). For a matrix

¹Technically, from the perspective of Lebesgue measure, we should use the more general class of

$A \in \mathbb{R}^{n \times n}$, we have the important relation $\text{vol}_n(AS) = |\det(A)|\text{vol}_n(S)$. In particular, we note that if A is orthogonal (i.e. $A^T A = I_n$), we have that $\text{vol}_n(AS) = \text{vol}_n(S)$. Lastly, for a scaling $t \geq 0$, this gives that $\text{vol}_n(tS) = t^n \text{vol}_n(S)$ (note that scaling is equivalent to the action of tI_n on S).

We define $\text{vol}_k(\cdot)$ on \mathbb{R}^n , $k \leq n$, as the normalized k dimensional Hausdorff measure on \mathbb{R}^n . In this thesis, we will only use vol_k to measure subsets of a k dimensional affine space $H = W + \mathbf{x} \subseteq \mathbb{R}^n$, where $W \subseteq \mathbb{R}^n$ is a k dimensional linear subspace and $\mathbf{x} \in \mathbb{R}^n$. In this setting vol_k is easy to understand. Let $P = (\mathbf{b}_1, \dots, \mathbf{b}_k) \in \mathbb{R}^{n \times k}$ denote a matrix whose columns form an orthonormal basis of W . Then for $S \subseteq H$, we have that $\text{vol}_k(S) = \text{vol}_k(P^T S)$. Since P^T maps S to \mathbb{R}^k , the second volume is simply corresponds to the standard k -dimensional Lebesgue measure on \mathbb{R}^k .

Integration. Let $(\Omega, \mathcal{A}, \mu)$ be a measure space. For a set $A \in \mathcal{A}$, we define $1_A : \Omega \rightarrow \mathbb{R}$, the indicator of A , by the relation $1_A(\mathbf{x}) = 1$ if $\mathbf{x} \in A$ and 0 otherwise.

A function $g : \Omega \rightarrow \mathbb{R}_+$ is simple if $g = \sum_{i=1}^k a_i 1_{A_i}$, where $A_i \in \mathcal{A}$ and $a_i \in \mathbb{R}_+$ for $i \in [k]$. We define the Lebesgue integral of g with respect to the measure μ as

$$\int g(\mathbf{x}) d\mu(\mathbf{x}) = \sum_{i=1}^k a_i \mu(A_i)$$

For a \mathcal{A} -measurable function $f : \Omega \rightarrow \mathbb{R}_+$, we define the Lebesgue integral of f as

$$\int f(\mathbf{x}) d\mu(\mathbf{x}) = \sup \left\{ \int g(\mathbf{x}) d\mu(\mathbf{x}) : g : \Omega \rightarrow \mathbb{R}_+ \text{ simple and } g \leq f \right\}$$

f is μ -integrable if $\int f d\mu < \infty$. If $f, g : \Omega \rightarrow \mathbb{R}_+$ are μ -integrable then $f + g$ is μ -integrable and

$$\int (f + g)(\mathbf{x}) d\mu(\mathbf{x}) = \int f(\mathbf{x}) d\mu(\mathbf{x}) + \int g(\mathbf{x}) d\mu(\mathbf{x})$$

Let $f_i : \Omega \rightarrow \mathbb{R}_+$, $i \in \mathbb{N}$, denote a monotonically non-decreasing sequence of functions,

Lebesgue measurable sets. However this generality will not be needed here.

i.e. $f_i \leq f_{i+1}$. Then the monotone convergence theorem states that

$$\lim_{i \rightarrow \infty} \int f_i(\mathbf{x}) d\mu(\mathbf{x}) = \int \lim_{i \rightarrow \infty} f_i(\mathbf{x}) d\mu(\mathbf{x})$$

For a \mathcal{A} -measurable function $f : \Omega \rightarrow \mathbb{R}$ (i.e. f is not necessarily non-negative), we say that f is μ -integrable if $|f|$ is μ -integrable. Here we define

$$\int f(\mathbf{x}) d\mu(\mathbf{x}) = \int f^+(\mathbf{x}) d\mu(\mathbf{x}) - \int f^-(\mathbf{x}) d\mu(\mathbf{x}),$$

where $f^+(\mathbf{x}) = \max\{0, f(\mathbf{x})\}$ and $f^- = \max\{0, -f(\mathbf{x})\}$ (both non-negative functions), noting that $f(\mathbf{x}) = f^+(\mathbf{x}) - f^-(\mathbf{x})$.

For convenience, for $A \in \mathcal{A}$, we write $\int_A f(\mathbf{x}) d\mu(\mathbf{x}) = \int f(\mathbf{x}) 1_A(\mathbf{x}) d\mu(\mathbf{x})$. If f is non-negative, then the function $\mu_f(A) = \int_A f(\mathbf{x}) d\mu(\mathbf{x})$ for $A \in \mathcal{A}$ defines a measure on \mathcal{A} .

For measurable functions $f, g : \Omega \rightarrow \mathbb{R}$, and a number $p \geq 1$, the following is the Minkowski inequality

$$\left(\int |f(\mathbf{x}) + g(\mathbf{x})|^p d\mu(\mathbf{x}) \right)^{\frac{1}{p}} \leq \left(\int |f(\mathbf{x})|^p d\mu(\mathbf{x}) \right)^{\frac{1}{p}} + \left(\int |g(\mathbf{x})|^p d\mu(\mathbf{x}) \right)^{\frac{1}{p}}$$

Let μ_1, μ_2 denote measures on the same space Ω and σ -algebra \mathcal{A} . We say that μ_2 is absolutely continuous with respect to μ_1 , if there exists a μ_1 -integrable function $f : \Omega \rightarrow \mathbb{R}_+$ such that for all $A \in \mathcal{A}$

$$\mu_2(A) = \int_A f(\mathbf{x}) d\mu_1(\mathbf{x})$$

Here we refer to f as the density of μ_2 with respect to μ_1 . In this thesis, we will mainly be interested in measures which are absolutely continuous with respect to the Lebesgue measure vol_n in \mathbb{R}^n . If a measure μ is absolutely continuous with respect to the Lebesgue measure, for $A \subseteq \mathbb{R}^n$ measurable, we may write

$$\mu(A) = \int_A f(\mathbf{x}) d\text{vol}_n(\mathbf{x})$$

for some measurable $f : \mathbb{R}^n \rightarrow \mathbb{R}_+$. For notational convenience we simply write $d\mathbf{x}$ for $d\text{vol}_n(\mathbf{x})$ when the context is clear, i.e. $\mu(A) = \int_A f(\mathbf{x}) d\mathbf{x}$.

Take a measurable function $f : \Omega \rightarrow \mathbb{R}^n$, which we write $(f_1, \dots, f_n)^T$ (i.e. a column vector). Here f is μ -integrable if each f_i , $i \in [n]$, is μ -integrable, where we define

$$\int f(\mathbf{x})d\mu(\mathbf{x}) = \left(\int f_1(\mathbf{x})d\mu(\mathbf{x}), \dots, \int f_n(\mathbf{x})d\mu(\mathbf{x}) \right)^T$$

Random Variables: Let $(\Omega, \mathcal{A}, \mu)$ denote a probability space (i.e. $\mu(\Omega) = 1$). A random variable is a measurable function $X : \Omega \rightarrow \mathbb{R}^n$. For a measurable $A \subseteq \mathbb{R}^n$, we denote the probability that $X \in A$ as

$$\Pr[X \in A] = \mu(\{\mathbf{w} \in \Omega : X(\mathbf{w}) \in A\})$$

We note that $P_X(A) = \Pr[X \in A]$ defines a probability measure on \mathbb{R}^n with σ -algebra \mathcal{B}^n . We call P_X the probability distribution of X on \mathbb{R}^n . We say that a random variable X is discrete if it takes on only countably many values, i.e. the set $X(\Omega)$ is countable. If X takes on a continuum of values (e.g. with support on $[0, 1]$), we denote X a continuous random variable.

For a measurable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we denote the expectation of f with respect to X as

$$\mathbb{E}[f(X)] \stackrel{\text{def}}{=} \int f(X(\mathbf{w}))d\mu(\mathbf{w}) = \int f(\mathbf{x})dP_X(\mathbf{x})$$

From the above, we see that to measure properties of the random variable X , the probability distribution P_X contains all the required information.

Given another random variable $Y : \Omega \rightarrow \mathbb{R}^n$, we note that $(X, Y) \in \mathbb{R}^{n \times 2}$ is again a random variable (by measurability of (X, Y)). For measurable $A \subseteq \mathbb{R}^{n \times 2}$, we denote the joint probability distribution of (X, Y) by

$$\Pr[(X, Y) \in A] \stackrel{\text{def}}{=} \mu(\{\mathbf{w} : (X(\mathbf{w}), Y(\mathbf{w})) \in A\})$$

The random variables X and Y are independent if for all $A, B \subseteq \mathbb{R}^n$ measurable

$$\Pr[X \in A, Y \in B] = \Pr[X \in A] \Pr[Y \in B].$$

Let P_X, P_Y denote probability distributions of X, Y . By the independence assumption on X and Y , we note that the joint probability distribution of (X, Y) is exactly the product measure $P_X \times P_Y$ (by uniqueness of the product measure). Hence for independent random variables, all the requisite information for the joint distribution is contained in the individual probability measures. Therefore we may define the joint distributions of such variables without reference to the ambient probability space.

In this thesis, we will often need to examine the distributions of sums of independent random variables. In particular, for X and Y independent random variables as above, we have that for $A \subseteq \mathbb{R}^n$ measurable the distribution of $X + Y$ satisfies

$$\Pr[X + Y \in A] = \int 1_A[\mathbf{x} + \mathbf{y}]d(P_X \times P_Y)(\mathbf{x}, \mathbf{y}) = \int \int 1_A[\mathbf{x} + \mathbf{y}]dP_X(\mathbf{x})dP_Y(\mathbf{y})$$

Here the probability distribution of $X + Y$ is known as the convolution of P_X and P_Y , which we write $P_X * P_Y$. If the probability distributions P_X, P_Y are absolutely continuous with respect to the Lebesgue measure then so is $P_X * P_Y$. Furthermore, if P_X, P_Y have associated densities $f, g : \mathbb{R}^n \rightarrow \mathbb{R}_+$ then the density of $P_X * P_Y$ is

$$(f * g)(\mathbf{x}) = \int f(\mathbf{x} - \mathbf{y})g(\mathbf{y})d\mathbf{y},$$

i.e. the convolution of f and g .

Let σ_1, σ_2 denote two probability measures on a domain Ω with σ -algebra \mathcal{A} . We define their *total variation* (or *statistical distance*) as

$$d_{\text{TV}}(\sigma_1, \sigma_2) = \sup_{A \in \mathcal{A}} |\sigma_1(A) - \sigma_2(A)|.$$

Gaussian Random Variables. For a measurable subset $A \subseteq \mathbb{R}^n$, we define the n dimensional Gaussian measure of A (parameterized by $s > 0$) as

$$\gamma_{n,s}(A) = \frac{1}{s^n} \int_A e^{-\pi \|\frac{\mathbf{x}}{s}\|_2^2} d\mathbf{x}.$$

From classical probability, $\gamma_{n,s}$ is a probability measure, i.e. $\gamma_{n,s}(\mathbb{R}^n) = 1$. Note that Gaussian measure is absolutely continuous with respect to the Lebesgue measure and

has density $s^{-n}e^{-\pi\|\frac{\mathbf{x}}{s}\|_2^2}$. We define the canonical (or standard) Gaussian measure on \mathbb{R}^n to be $\gamma_n \stackrel{\text{def}}{=} \gamma_{n,\sqrt{2\pi}}$.

Define the Gaussian probability distribution on \mathbb{R}^n with parameter $s > 0$, center $\mathbf{c} \in \mathbb{R}^n$, as

$$D_{n,s,\mathbf{c}}(A) = \gamma_{n,s}(A - \mathbf{c}) \text{ for } A \subseteq \mathbb{R}^n \text{ measurable.}$$

Here we define $D_{n,s} \stackrel{\text{def}}{=} D_{n,s,\mathbf{0}}$. When the context is clear, we drop the n in the notation, and write D_s and $D_{s,\mathbf{c}}$ for $D_{n,s}$ and $D_{n,s,\mathbf{c}}$

We say that $X \in \mathbb{R}^n$ is a Gaussian random vector with distribution $D_{n,s,\mathbf{c}}$ if $\Pr[X \in A] = D_{n,s,\mathbf{c}}(A)$, for all $A \subseteq \mathbb{R}^n$ measurable. From classical probability, we have that $\mathbb{E}[X] = \mathbf{c}$ and that for $\mathbf{v} \in \mathbb{R}^n$, $\mathbb{E}[\langle X - \mathbf{c}, \mathbf{v} \rangle^2] = \|\mathbf{v}\|^2 \left(\frac{s}{\sqrt{2\pi}}\right)^2$. We say that $X \in \mathbb{R}^n$ is standard Gaussian if it has distribution $D_{n,\sqrt{2\pi}}$. For $X \in \mathbb{R}^n$ standard Gaussian, $\forall \mathbf{v} \in \mathbb{R}^n$, we have that $\mathbb{E}[\langle X, \mathbf{v} \rangle^2] = \|\mathbf{v}\|_2^2$, or equivalently $\mathbb{E}[XX^T] = I_n$. Furthermore, we see that $\mathbb{E}[\|X\|_2^2] = \sum_{i=1}^n \mathbb{E}[X_i^2] = n$.

If $X \in \mathbb{R}^n$ is a Gaussian random vector with distribution D_{n,s_1,\mathbf{c}_1} , then the scaling tX , $t > 0$, is distributed as $D_{n,ts_1,t\mathbf{c}_1}$. If $Y \in \mathbb{R}^n$ is a Gaussian random vector independent from X with distribution D_{n,s_2,\mathbf{c}_2} , then the sum $X + Y$ is again Gaussian and is distributed as $D_{n,\sqrt{s_1^2+s_2^2},\mathbf{c}_1+\mathbf{c}_2}$.

2.2 Convexity

For detailed exposition of convex analysis, the structure of convex sets, and polyhedral theory, the reader may consult [110, 117, 120].

Convex Sets. A set $K \subseteq \mathbb{R}^n$ is convex if for all $\mathbf{x}, \mathbf{y} \in K$, the line segment $[\mathbf{x}, \mathbf{y}] \in K$. $K \subseteq \mathbb{R}^n$ is a convex body if K is convex, compact and full-dimensional. K is $\mathbf{0}$ -centered if $\mathbf{0} \in \text{int}(K)$. K is centrally symmetric if $K = -K$.

We state the following simple lemma, which will be useful in the segway.

Lemma 2.2.1.

- (1) $K \subseteq \mathbb{R}^n$ is a convex set if and only if $\forall a, b \geq 0, aK + bK = (a + b)K$.
- (2) If $K \subseteq \mathbb{R}^n$ is a convex set with $\mathbf{0} \in K$, then for all $0 \leq t \leq s, tK \subseteq sK$.
- (3) If $K \subseteq \mathbb{R}^n$ is a symmetric convex body, then K is a $\mathbf{0}$ -centered convex body.

Separation. A halfspace is a subset of \mathbb{R}^n defined by a single linear inequality. For $\mathbf{v} \in \mathbb{R}^n, a \in \mathbb{R}$, we define the halfspace $H_{\mathbf{v},a}^{\leq(\geq)} = \{\mathbf{x} : \langle \mathbf{v}, \mathbf{x} \rangle \leq(\geq)a\}$. For $\mathbf{v} \neq \mathbf{0}$, we define the hyperplane $H_{\mathbf{v},a} = \{\mathbf{x} : \langle \mathbf{v}, \mathbf{x} \rangle = a\}$, where we note that $H_{\mathbf{v},a} = \partial H_{\mathbf{v},a}^{\leq(\geq)}$. Henceforth, we shall write H to denote a generic hyperplane and H^{\leq} for a generic halfspace.

One of the most important concepts in convexity is that of a separator. Given two convex sets $K, C \subseteq \mathbb{R}^n$, we say that a hyperplane H separates K and C , if $\forall \mathbf{x} \in K, \mathbf{y} \in C$, the line $[\mathbf{x}, \mathbf{y}] \cap H \neq \emptyset$ and $K \cup C \subsetneq H$. Furthermore, we say that H strictly separates K and C (or H is a strict separator) if $K \cap H = C \cap H = \emptyset$. Similarly, we say that a halfspace H^{\leq} separates (strictly) separates K and C if the hyperplane ∂H^{\leq} (strictly) separates K and C .

The first fundamental theorem in convexity states that disjoint convex sets can always be separated:

Theorem 2.2.2 (Separator Theorem). *Let $K, C \subseteq \mathbb{R}^n$ denote non-empty convex sets. Then the following holds:*

- (1) *If $\text{relint}(K) \cap \text{relint}(C) = \emptyset$, there exists a hyperplane H separating K and C .*
- (2) *If $K \cap C = \emptyset$, K is compact and C is closed, there exists a hyperplane H which strictly separates K and C .*

From the separator theorem, we immediately get the following corollary, which tells us a first structural characterization of convex sets:

Theorem 2.2.3. *$K \subseteq \mathbb{R}^n$ is a closed convex set if and only if K can be expressed as an intersection of halfspaces (possibly infinite).*

A fundamental class of convex sets is the class of polyhedra: the class of bodies which can be defined as a finite intersection of halfspaces. More precisely, $P \subseteq \mathbb{R}^n$ is a polyhedron if P can be written as $\{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$ where $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$ (for vectors, $\mathbf{a} \leq \mathbf{b}$ if $\mathbf{a}_i \leq \mathbf{b}_i$ for all i). The following fundamental result, known as Farkas Lemma, tells us precisely when a linear inequality is valid for a polyhedron:

Lemma 2.2.4 (Farkas Lemma). *Let $P = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$, where $A \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$. Then for $\mathbf{v} \in \mathbb{R}^n$, $a \in \mathbb{R}$, we have that*

$$P \subseteq H_{\mathbf{v},a}^{\leq} \quad \Leftrightarrow \quad \exists \mathbf{y} \in \mathbb{R}_+^m \text{ such that } A^T \mathbf{y} = \mathbf{v} \text{ and } \langle \mathbf{y}, \mathbf{b} \rangle \leq a.$$

Faces. Let $K \subseteq \mathbb{R}^n$ denote a convex set. We say that $F \subseteq K$ is a face of K if for every line segment $[x, y] \subseteq K$, $[x, y] \cap F \neq \emptyset \Rightarrow [x, y] \subseteq F$. A face F of K is proper if $F \neq K$. A facet $F \subseteq K$ is $d-1$ dimensional face of K , where $d = \dim(K)$. An extreme point (or vertex) of K is a 0-dimensional face of K . We denote the set of extreme points of K as $\text{ext}(K)$. We say that a hyperplane H is a supporting hyperplane of K , if $H \cap K$ is a non-empty face of K . Similarly, H^{\leq} is a supporting halfspace of (or bounds a face of) K if $K \subseteq H^{\leq}$ and $K \cap \partial H^{\leq}$ is a non-empty face of K . A halfspace H^{\leq} is facet defining for K if $K \cap \partial H^{\leq}$ is a facet of K . A face $F \subseteq K$ is supported by a (halfspace) hyperplane (H^{\leq}) H , if (H^{\leq}) H is a supporting (halfspace) hyperplane for K and $F \subseteq (H^{\leq})H$. Furthermore, F is an exposed face if there exists a hyperplane H such that $H \cap K = F$. If K is a polyhedron, then one can show that all of K 's faces are exposed.

Lastly, we note that Theorem 2.2.3 can be strengthened by saying that a closed convex set K is the intersection of its supporting halfspaces. Furthermore for a full dimensional polyhedron P , this can be strengthened further by saying that P is the intersection of all its facet defining halfspaces.

Algebra of Convex Sets. We give some constructions of convex sets and their properties. Let $K, C \subseteq \mathbb{R}^n$ be convex sets. Then the sets $K \cap C$, $K - C$ and $K + C$ are all convex. For linear transformations $T_1 : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $T_2 : \mathbb{R}^m \rightarrow \mathbb{R}^n$, the image $T_1(K)$ and inverse image $T_2^{-1}(K)$ are convex. For a subset $A \subseteq \mathbb{R}^n$, define the convex hull $\text{conv}(A)$ as the smallest convex set (by inclusion) containing A . For K and C as before, we have that $\text{conv}\{K, C\} \stackrel{\text{def}}{=} \text{conv}\{K \cup C\} = \bigcup_{0 \leq \lambda \leq 1} \lambda K + (1 - \lambda)C$.

A fundamental construct in convex analysis is the polar of a convex set. Let K be a closed convex set such that $\mathbf{0} \in K$. We define the *polar* of K^* as

$$K^* = \{\mathbf{y} \in \text{span}(K) : \forall \mathbf{x} \in K, \langle \mathbf{x}, \mathbf{y} \rangle \leq 1\}$$

Here note that K^* is convex and that $\mathbf{0} \in K^*$. A fundamental theorem in convex analysis (which follows from the separator theorem) is that $(K^*)^* = K$. Furthermore, for a linear subspace $W \subseteq \mathbb{R}^n$, we have the relation $(K \cap W)^* = \pi_W(K^*)$.

Functions over Convex Sets. A function $f : K \rightarrow \bar{\mathbb{R}}_+$ is convex, if the domain K of f is a convex set, and $\forall \mathbf{x}, \mathbf{y} \in K$, and $\alpha \in [0, 1]$, f satisfies

$$f(\alpha \mathbf{x} + (1 - \alpha)\mathbf{y}) \leq \alpha f(\mathbf{x}) + (1 - \alpha)f(\mathbf{y}) \quad (2.2.1)$$

From the definition, one can check that the level sets of f , i.e. $\{\mathbf{x} \in K : f(\mathbf{x}) \leq a\}$ for $a \in \mathbb{R}^n$, are all convex. Note that if $f, g : K \rightarrow \bar{\mathbb{R}}$ are convex functions, then $f + g$ and $\max\{f, g\}$ are also convex. Lastly, we say that a function $f : K \rightarrow \mathbb{R}$ is concave if $-f$ is convex.

A subgradient $\mathbf{v} \in \mathbb{R}^n$ of f at $\mathbf{x} \in K$, written $\mathbf{v} \in \partial f(\mathbf{x})$, is a vector satisfying

$$f(\mathbf{z}) \geq f(\mathbf{x}) + \langle \mathbf{v}, \mathbf{z} - \mathbf{x} \rangle \quad \forall \mathbf{z} \in K$$

A classical fact from convex analysis is that a convex function admits subgradients at all points in its domain.

Let $X : \Omega \rightarrow K$ be a random variable, and let $f : K \rightarrow \mathbb{R}$ be a convex function. A fundamental inequality we shall use is Jensen's inequality, which states that

$$f(\mathbb{E}[X]) \leq \mathbb{E}[f(X)] \quad (2.2.2)$$

Implicit in the above inequality is that the vector $\mathbb{E}[X] \in K$. This follows by convexity of K and the fact that $\mathbb{E}[X]$ is an average of points in K .

We define the support function of a convex set $K \subseteq \mathbb{R}^n$ by

$$h_K(\mathbf{v}) = \sup_{\mathbf{x} \in K} \langle \mathbf{v}, \mathbf{x} \rangle, \text{ for } \mathbf{v} \in \mathbb{R}^n.$$

Since h_K is the supremum of linear functions (which are convex), we get that $h_K : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}_+$ is a convex function.

We define width_K , the width functional of K by

$$\text{width}_K(\mathbf{y}) = \sup_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle - \inf_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle = h_K(\mathbf{y}) + h_K(-\mathbf{y}).$$

If K is a convex body, then $\text{width}_K : \mathbb{R}^n \rightarrow \mathbb{R}_+$ in fact defines a norm on \mathbb{R}^n . We will use this in the sequel.

Let $K \subseteq \mathbb{R}^n$ be convex set. We define the gauge function, or Minkowski functional, $\|\cdot\|_K$ of K by $\|\mathbf{x}\|_K = \inf\{s \geq 0 : \mathbf{x} \in sK\}$, $\mathbf{x} \in \mathbb{R}^n$.

Let $C_1, C_2, K \subseteq \mathbb{R}^n$ be convex sets satisfying the containment relation $C_1 \subseteq K \subseteq C_2$. Then by the definition of the gauge function, it is easy to see that for $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_{C_2} \leq \|\mathbf{x}\|_K \leq \|\mathbf{x}\|_{C_1}$. Furthermore, If for $0 < a \leq b$, we have that containment relations $aC_1 \subseteq K \subseteq bC_1$, then for all $\mathbf{x} \in \mathbb{R}^n$, we have that $\frac{1}{b}\|\mathbf{x}\|_{C_1} \leq \|\mathbf{x}\|_K \leq \frac{1}{a}\|\mathbf{x}\|_{C_1}$.

The next lemma shows that convex bodies are in a sense equivalent to norms. This fact will be used continuously throughout the thesis.

Lemma 2.2.5. *Take $S \subseteq \mathbb{R}^n$. Then the following holds:*

- (1) $\|\cdot\|_S$ is an asymmetric norm satisfying $S = \{\mathbf{x} : \|\mathbf{x}\|_S \leq 1\}$ if and only if S is a $\mathbf{0}$ -centered convex body.

(2) $\|\cdot\|_S$ is an symmetric norm satisfying $S = \{\mathbf{x} : \|\mathbf{x}\|_S \leq 1\}$ if and only if S is a symmetric convex body.

The last lemma of this section establishes certain dualities between the support function and gauge function of a convex body.

Lemma 2.2.6. *Let $K \subseteq \mathbb{R}^n$ be $\mathbf{0}$ -centered convex body. Then the following holds:*

(1) For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we have that $\langle \mathbf{x}, \mathbf{y} \rangle \leq \|\mathbf{x}\|_K \|\mathbf{y}\|_{K^*}$.

(2) For $\mathbf{v} \in \mathbb{R}^n$, $\|\mathbf{x}\|_K = \inf\{s \geq 0 : \mathbf{x} \in sK\} = \sup_{\mathbf{z} \in K^*} \langle \mathbf{x}, \mathbf{z} \rangle = h_{K^*}(\mathbf{x})$.

Furthermore, for a linear subspace $W \subseteq \mathbb{R}^n$,

$$\|\mathbf{x}\|_{\pi_W(K)} = \inf_{\mathbf{z} \in W^\perp} \|\mathbf{x} + \mathbf{z}\|_K = \sup_{\mathbf{z} \in K^* \cap W} \langle \mathbf{x}, \mathbf{z} \rangle = h_{K^* \cap W}(\mathbf{x})$$

2.3 Convex Geometry

For a treatment of the fundamental results in Convex Geometry, the reader may consult [117, 9, 53].

Definition 2.3.1 (Distances Measures for Convex Bodies). For convex bodies $K, C \subseteq \mathbb{R}^n$ we define their geometric distance as

$$d(K, C) = \inf\left\{\frac{b}{a} : a, b > 0, \mathbf{x} \in K, \mathbf{y} \in C, a(K - \mathbf{x}) \subseteq C - \mathbf{y} \subseteq b(K - \mathbf{x})\right\}$$

We define their Banach-Mazur distance as

$$d_{BM}(K, C) = \inf\{s : \mathbf{x} \in K, \mathbf{y} \in C, T : \mathbb{R}^n \rightarrow \mathbb{R}^n \text{ non-singular linear transformation,}$$

$$T(K - \mathbf{x}) \subseteq C - \mathbf{y} \subseteq sT(K - \mathbf{x})\}$$

A fundamental theorem in Convex Geometry is John's theorem, which bounds the distance of any convex body the euclidean ball.

Theorem 2.3.2 (John's Theorem). *If $K \subseteq \mathbb{R}^n$ is a convex body, then $d_{BM}(K, B_2^n) \leq n$. Furthermore, if K is centrally symmetric then $d_{BM}(K, B_2^n) \leq \sqrt{n}$.*

From the measure theoretic perspective, the first fundamental inequality is the Brunn-Minkowski inequality, which shows the volume is in a sense concave.

Theorem 2.3.3 (Brunn Minkowski). *Let $A, B \subseteq \mathbb{R}^n$ be non-empty compact sets. Then*

$$\text{vol}_n(A)^{\frac{1}{n}} + \text{vol}_n(B)^{\frac{1}{n}} \leq \text{vol}(A + B)^{\frac{1}{n}}$$

2.3.1 Logconcave Functions and Convex Bodies

A very useful concept in convex geometry is that of the logconcave function, which in a sense generalizes both convex functions and convex bodies. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is logconcave if for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, and $0 \leq \alpha \leq 1$, we have that

$$f(\alpha \mathbf{x} + (1 - \alpha)\mathbf{y}) \geq f(\mathbf{x})^\alpha f(\mathbf{y})^{1-\alpha}$$

Equivalently, f is logconcave if and only if $\log f$ is concave. Some examples of logconcave functions are indicator functions of convex bodies, the gaussian density function, concave functions, and much more.

For a logconcave function f on \mathbb{R}^n such that $0 < \int_{\mathbb{R}^n} f(\mathbf{x}) \, d\mathbf{x} < \infty$, we define the associated probability measure (distribution) π_f , where for measurable $A \subseteq \mathbb{R}^n$, we have

$$\pi_f(A) = \frac{\int_A f(\mathbf{x}) \, d\mathbf{x}}{\int_{\mathbb{R}^n} f(\mathbf{x}) \, d\mathbf{x}}.$$

We define the *centroid* (or barycenter) and *covariance* matrix of f as

$$\mathbf{b}(f) = \frac{\int_{\mathbb{R}^n} \mathbf{x} f(\mathbf{x}) \, d\mathbf{x}}{\int_{\mathbb{R}^n} f(\mathbf{x}) \, d\mathbf{x}} \quad \text{cov}(f)_{ij} = \frac{\int_{\mathbb{R}^n} (\mathbf{x}_i - \mathbf{b}(f)_i)(\mathbf{x}_j - \mathbf{b}(f)_j) f(\mathbf{x}) \, d\mathbf{x}}{\int_{\mathbb{R}^n} f(\mathbf{x}) \, d\mathbf{x}} \quad 1 \leq i, j \leq n$$

The matrix $\text{cov}(f)$ is positive semi-definite and symmetric. We say that f is isotropic, or in isotropic position, if $\mathbf{b}(f) = \mathbf{0}$ and $\text{cov}(f)$ is the identity matrix. Define the *inertial ellipsoid* of f as

$$E_f = E(\text{cov}(f)^{-1}) = \{\mathbf{x} : \mathbf{x}^t \text{cov}(f)^{-1} \mathbf{x} \leq 1\}$$

The *isotropic constant* of f is defined as

$$L_f = \left(\sup_{\mathbf{x} \in \mathbb{R}^n} \frac{f(\mathbf{x})}{\int_{\mathbb{R}^n} f(\mathbf{x}) d\mathbf{x}} \right)^{\frac{1}{n}} \cdot \det(\text{cov}(f))^{\frac{1}{2n}}.$$

For a convex body K , let π_K denote the uniform measure (distribution) over K . Let f_K denote the associated density, i.e.,

$$f_K(\mathbf{x}) = \frac{1}{\text{vol}_n(K)} I[\mathbf{x} \in K],$$

We extend all the above notions to convex bodies by defining $\text{cov}(K) \stackrel{\text{def}}{=} \text{cov}(f_K)$, $\mathbf{b}(K) \stackrel{\text{def}}{=} \mathbf{b}(f_K)$, $L_K \stackrel{\text{def}}{=} L_{f_K}$, etc. We say that K is in isotropic position if $\mathbf{b}(K) = \mathbf{0}$ and $\text{cov}(K)$ is the identity (a different normalization is sometimes used in asymptotic convex geometry, namely, $\mathbf{b}(K) = \mathbf{0}$, $\text{vol}_n(K) = 1$, and $\text{cov}(K)$ is constant diagonal).

A major open conjecture in convex geometry is the following:

Conjecture 2.3.4 (Slicing Conjecture [20]). There exists an absolute constant $C > 0$, such that $L_K \leq C$ for all $n \geq 1$ for any convex body $K \subseteq \mathbb{R}^n$.

The original bound computed by Bourgain [20] was $L_K \leq Cn^{1/4} \log n$, $C > 0$ an absolute constant. This has since been improved by Klartag [78] to $L_K = Cn^{1/4}$, $C > 0$ an absolute constant. In addition, the conjecture has been verified for many classes of bodies such as 1-unconditional bodies, zonoids, duals of zonoids, etc. We note that associated conjecture for logconcave functions (i.e. $L_f \leq C$ for a logconcave function) is equivalent to the slicing conjecture up to a universal constant [8].

2.3.2 Geometric Inequalities

The following gives bounds on how well the inertial ellipsoid approximates a convex body. The estimates below are from [68]:

Theorem 2.3.5. For a convex body $K \subseteq \mathbb{R}^n$, the inertial ellipsoid E_K satisfies

$$\sqrt{\frac{n+2}{n}} \cdot E_K \subseteq K - \mathbf{b}(K) \subseteq \sqrt{n(n+2)} \cdot E_K$$

The above containment relationship was shown in [94] for centrally symmetric bodies (with better bounds), and by [124] for general bodies with suboptimal constants.

The next theorem gives estimates on the volume product, a fundamental quantity in Convex Geometry. The upper bound for centrally symmetric bodies follows from the work of Blaschke [17], and for general bodies by Santaló [116]. The lower bound was first established by Bourgain and Milman [21], and was recently refined by Kuperberg [79], as well as by Nazarov [98], where Kuperberg achieves the best constants. Finding the exact minimizer of the volume product is a major open problem in Convex Geometry.

Theorem 2.3.6. *Let K be a convex body in \mathbb{R}^n . Then we have*

$$\text{vol}_n(B_2^n)^2 \geq \inf_{\mathbf{x} \in K} \text{vol}_n(K - \mathbf{x}) \text{vol}_n((K - \mathbf{x})^*) \geq \left(\frac{\pi e(1 + o(1))}{2n} \right)^n.$$

If K is centrally symmetric, then

$$\text{vol}_n(B_2^n)^2 \geq \text{vol}_n(K) \text{vol}_n(K^*) \geq \left(\frac{\pi e(1 + o(1))}{n} \right)^n.$$

In both cases, the upper bounds are equalities if and only if K is an ellipsoid.

We remark that the upper and lower bounds match within a 4^n factor (2^n for symmetric bodies) since $\text{vol}_n(B_2^n)^2 = \left(\frac{2\pi e(1+o(1))}{n} \right)^n$.

The next theorem gives useful volume estimates for some basic operations on a convex body. The first estimate is due to Rogers and Shepard [111], and the second is due Milman and Pajor [95]:

Theorem 2.3.7. *Let $K \subseteq \mathbb{R}^n$ be a convex body. Then*

$$\text{vol}_n(K - K) \leq \binom{2n}{n} \text{vol}_n(K) \leq 4^n \text{vol}_n(K).$$

If $\mathbf{b}(K) = \mathbf{0}$, i.e., the centroid of K is at the origin, then

$$\text{vol}_n(K) \leq 2^n \text{vol}_n(K \cap -K).$$

Definition 2.3.8 (Covering Numbers). For convex bodies $K, T \subseteq \mathbb{R}^n$, we define

$$N(K, T) = \min\{|\Lambda| : \Lambda \subseteq \mathbb{R}^n, K \subseteq \Lambda + T\},$$

to be the minimum number of translates of T needed to cover K .

We relate some well-known covering estimates.

Lemma 2.3.9. *Let $K, T \subseteq \mathbb{R}^n$ be convex bodies. Then*

$$N(K, T) \leq 6^n \inf_{c \in \mathbb{R}^n} \frac{\text{vol}_n(K)}{\text{vol}_n(K \cap (T + c))} \quad \text{and} \quad \frac{\text{vol}_n(K + T)}{\text{vol}_n(T)} \leq 2^n N(K, T).$$

If T is centrally symmetric, then

$$N(K, T) \leq \frac{\text{vol}_n(K + T/2)}{\text{vol}_n(T/2)}.$$

If $K \cap T$ are centrally symmetric, then

$$N(K, T) \leq 3^n \frac{\text{vol}_n(K)}{\text{vol}_n(K \cap T)}.$$

Proof. Let us first examine the case where T is centrally symmetric, where we wish to show that

$$N(K, T) \leq \frac{\text{vol}_n(K + T/2)}{\text{vol}_n(T/2)} \tag{2.3.1}$$

Let $\Lambda \subseteq K$ be a maximal subset of K such that for $\mathbf{x}_1, \mathbf{x}_2 \in \Lambda$, $\mathbf{x}_1 \neq \mathbf{x}_2$, $\mathbf{x}_1 + T/2 \cap \mathbf{x}_2 + T/2 = \emptyset$.

Claim 1: $K \subseteq \cup_{\mathbf{x} \in \Lambda} \mathbf{x} + T$.

Take $\mathbf{y} \in K$. By maximality of Λ , there exists $\mathbf{x} \in \Lambda$ such that

$$\mathbf{y} + T/2 \cap \mathbf{x} + T/2 \neq \emptyset \quad \Rightarrow \quad \mathbf{y} \in \mathbf{x} + T/2 - T/2 \quad \Rightarrow \quad \mathbf{y} \in \mathbf{x} + T$$

where the last equality follows since T is centrally symmetric. The claim thus follows.

Claim 2: $|\Lambda| \leq \frac{\text{vol}_n(K + T/2)}{\text{vol}_n(T/2)}$.

For $\mathbf{x} \in \Lambda$, note that since $\mathbf{x} \in K$, we have that $\mathbf{x} + T/2 \subseteq K + T/2$. Therefore $\Lambda + T/2 \subseteq K + T/2$. Since the sets $\mathbf{x} + T/2$, $\mathbf{x} \in \Lambda$, are disjoint, we have that

$$\text{vol}_n(K + T/2) \geq \text{vol}_n(\Lambda + T/2) = |\Lambda| \text{vol}_n(T/2) \quad (2.3.2)$$

as needed.

If $K \cap T$ is centrally symmetric, then by the estimate in (2.3.1) we get that

$$\begin{aligned} N(K, T) &\leq N(K, T \cap K) \leq \frac{\text{vol}_n(K + \frac{1}{2}(T \cap K))}{\text{vol}_n(\frac{1}{2}(T \cap K))} \\ &\leq \frac{\text{vol}_n(\frac{3}{2}K)}{\text{vol}_n(\frac{1}{2}(T \cap K))} = 3^n \frac{\text{vol}_n(K)}{\text{vol}_n(T \cap K)} \end{aligned} \quad (2.3.3)$$

as needed.

Now we examine the case where neither K nor T is necessarily symmetric. Since the covering estimate is shift invariant, we may assume that K and T have been shifted such that $\text{vol}_n(K \cap T)$ is maximized, and that the centroid of $K \cap T$ is at $\mathbf{0}$. Let $S = (K \cap T) \cap -(K \cap T)$. By Theorem 2.3.7 we have that $\text{vol}_n(S) \geq 2^{-n} \text{vol}_n(K \cap T)$. Note that S is a centrally symmetric convex body. Hence by identical reasoning as in (2.3.3) we get that

$$N(K, T) \leq 3^n \frac{\text{vol}_n(K)}{\text{vol}_n(S)} \leq 6^n \frac{\text{vol}_n(K)}{\text{vol}_n(K \cap T)}$$

as needed.

Lastly, pick any $\Lambda \subseteq \mathbb{R}^n$ such that $K \subseteq \Lambda + T$ and $|\Lambda| = N(K, T)$. Now we see that

$$\text{vol}_n(K + T) \leq \text{vol}_n((\Lambda + T) + T) = \text{vol}_n(\Lambda + 2T) \leq |\Lambda| \text{vol}_n(2T) = 2^n \text{vol}_n(T) N(K, T)$$

as needed. □

2.4 Lattices

For a thorough exposition on the fundamental properties of lattices, the reader may consult [25, 59].

A k -dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$ is a discrete subgroup under addition. It can be written as

$$\mathcal{L} = \left\{ \sum_{i=1}^k z_i \mathbf{b}_i : z_i \in \mathbb{Z}, i \in [k] \right\} \quad (2.4.1)$$

for some (not necessarily unique) *basis* $B = (\mathbf{b}_1, \dots, \mathbf{b}_k) \in \mathbb{R}^{n \times k}$ of $k \leq n$ linearly independent vectors in \mathbb{R}^n . In this thesis, we will interchangeably refer to the matrix B and its column vectors $\mathbf{b}_1, \dots, \mathbf{b}_k$ as a basis for \mathcal{L} .

A matrix $U \in \mathbb{R}^{k \times k}$ is unimodular if $U \in \mathbb{Z}^{k \times k}$ and $U^{-1} \in \mathbb{Z}^{k \times k}$. Equivalently by Cramer's rule (the matrix inversion formula), U is unimodular if and only if $U \in \mathbb{Z}^{k \times k}$ and $\det(U) = \pm 1$. Two bases $B_1, B_2 \in \mathbb{R}^{n \times k}$ generate the same lattice if and only if $B_1 = B_2 U$, for a $k \times k$ unimodular matrix U .

Let \mathcal{L} be k -dimensional lattices generated by $B = (\mathbf{b}_1, \dots, \mathbf{b}_k)$. The determinant of \mathcal{L} is defined as

$$\det(\mathcal{L}) = \sqrt{\det(\text{Gram}(\mathbf{b}_1, \dots, \mathbf{b}_k))} \stackrel{\text{def}}{=} \sqrt{\det(B^T B)}. \quad (2.4.2)$$

The determinant is a lattice invariant, i.e. it does not depend on the choice of basis for \mathcal{L} . To see this, for a unimodular matrix $U \in \mathbb{Z}^{k \times k}$, note that

$$\begin{aligned} \det((BU)^T (BU)) &= \det(U^T B^T B U) = \det(U^T) \det(B^T B) \det(U) \\ &= \det(U)^2 \det(B^T B) = \det(B^T B). \end{aligned}$$

Since every basis of \mathcal{L} can be expressed as BU , for some unimodular U , the above equation proves the claimed invariance.

The *dual lattice* \mathcal{L}^* of \mathcal{L} is defined as

$$\mathcal{L}^* = \{y \in \text{span}(\mathcal{L}) : \forall \mathbf{x} \in \mathcal{L}, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}. \quad (2.4.3)$$

A basis matrix $B^* \in \mathbb{R}^{n \times k}$ for the dual lattice (i.e. the columns form a basis of \mathcal{L}^*) can be obtained from a basis B of \mathcal{L} , by letting $B^* = B(B^T B)^{-1}$. We note that if \mathcal{L} is n dimensional, i.e. B forms a basis of \mathbb{R}^n , then the expression for B^* simplifies to B^{-T} .

Two fundamental relations between \mathcal{L} and \mathcal{L}^* are that $(\mathcal{L}^*)^* = \mathcal{L}$ and that $\det(\mathcal{L}) \det(\mathcal{L}^*) = 1$.

A linear subspace $W \subseteq \mathbb{R}^n$ is a lattice subspace of \mathcal{L} if \mathcal{L} contains a basis of W , or equivalently $\dim(\mathcal{L} \cap W) = \dim(W)$. The following lemma, tells us which some basic operations which preserve lattice structure, and well as some duality relations.

Lemma 2.4.1. *Let \mathcal{L} denote an n -dimensional lattice. The following holds:*

- (1) *If W is a lattice subspace of \mathcal{L} , then $\mathcal{L} \cap W$ is a lattice and $(\mathcal{L} \cap W)^* = \pi_W(\mathcal{L}^*)$.*
- (2) *For a linear map $T : \mathbb{R}^n \rightarrow W$, $W \subseteq \mathbb{R}^m$ a linear subspace, $T\mathcal{L}$ is a lattice if and only if $\text{Kern}(T)$ is a lattice subspace of \mathcal{L} .*
- (3) *For $W \subseteq \mathbb{R}^n$ a linear subspace, then $\pi_W(\mathcal{L})$ is a lattice if and only if W is a lattice subspace of \mathcal{L}^* .*

Let \mathcal{L} be an n -dimensional lattice. We define a coset of \mathcal{L} to be a set of the form $\mathcal{L} + \mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^n$. We define the equivalence relation $\mathbf{x} \equiv \mathbf{y} \pmod{\mathcal{L}} \Leftrightarrow \mathbf{x} - \mathbf{y} \in \mathcal{L}$. For a set $S \subseteq \mathbb{R}^n$ we write $S \pmod{\mathcal{L}}$ to denote set of equivalences classes represented in S . Here $\mathbb{R}^n \pmod{\mathcal{L}}$ corresponds to the set of all cosets of \mathcal{L} . Furthermore $\mathbb{R}^n \pmod{\mathcal{L}}$ forms a group under addition. A simple yet important lemma, tells us the structure of certain subgroups of $\mathbb{R}^n \pmod{\mathcal{L}}$.

Lemma 2.4.2. *Let \mathcal{L} denote an n dimensional lattice, and let $m \geq 1$ be an positive integer. Then the group $(\mathcal{L}/m \pmod{\mathcal{L}}, +)$ is isomorphic to \mathbb{Z}_m^n . In particular, $|\mathcal{L}/m \pmod{\mathcal{L}}| = m^n$.*

Let $K \subseteq \mathbb{R}^n$ be a convex body. The *covering radius* of \mathcal{L} with respect to K is $\mu(K, \mathcal{L}) = \inf\{s \geq 0 : \mathcal{L} + sK = \mathbb{R}^n\}$. Since $\mathbb{R}^n + \mathbf{t} = \mathbb{R}^n$ for any $\mathbf{t} \in \mathbb{R}^n$, we have that $\mu(K + \mathbf{t}, \mathcal{L}) = \mu(K, \mathcal{L})$. Furthermore, since $\mathcal{L} = -\mathcal{L}$ we also have that $\mu(-K, \mathcal{L}) = \mu(K, \mathcal{L})$.

Let $K \subseteq \mathbb{R}^n$ be a $\mathbf{0}$ -centered convex body. The length of the shortest non-zero vector (or *minimum distance*) of \mathcal{L} with respect to K is $\lambda_1(K, \mathcal{L}) = \min_{\mathbf{y} \in \mathcal{L} \setminus \{0\}} \|\mathbf{y}\|_K$.

For a point $\mathbf{x} \in \mathbb{R}^n$, we define the distance of \mathbf{x} to \mathcal{L} with respect to K as $d_K(\mathcal{L}, \mathbf{x}) = \inf_{\mathbf{y} \in \mathcal{L}} \|\mathbf{y} - \mathbf{x}\|_K$.

We define the i^{th} *minimum* of \mathcal{L} , $i \in [n]$, with respect to K

$$\lambda_i(K, \mathcal{L}) = \inf\{r \geq 0 : \dim(\text{span}(rK \cap \mathcal{L})) \geq i\}$$

For notational simplicity, for the ℓ_2 norm we write $\lambda_i(\mathcal{L}) = \lambda_i(B_2^n, \mathcal{L})$ and $\mu(\mathcal{L}) = \mu(B_2^n, \mathcal{L})$.

We note that all of the above concepts easily generalize to lower dimensional lattices.

2.4.1 Packing, Covering and Tiling

Let $F \subseteq \mathbb{R}^n$ be measurable set. We define F to be

- (1) \mathcal{L} -packing if $\forall \mathbf{x}, \mathbf{y} \in \mathcal{L}, \mathbf{x} \neq \mathbf{y}, (\mathbf{x} + F) \cap (\mathbf{y} + F) = \emptyset$
- (2) \mathcal{L} -covering if $\mathcal{L} + F = \mathbb{R}^n$.
- (3) \mathcal{L} -tiling (or a fundamental domain for \mathcal{L}) if F is both \mathcal{L} -packing and \mathcal{L} -covering.

We derive the following simple equivalence. F is \mathcal{L} -(packing, covering, tiling) if

$$\forall \mathbf{x} \in \mathbb{R}^n, \quad |(\mathcal{L} + \mathbf{x}) \cap F| \quad (\leq, \geq, =) \quad 1. \quad (2.4.4)$$

Assume F is \mathcal{L} -packing. Take $\mathbf{x} \in \mathbb{R}^n$. If $|(\mathcal{L} + \mathbf{x}) \cap F| \geq 2$, then we can pick distinct $\mathbf{w}, \mathbf{z} \in F$ such that $\mathbf{w}, \mathbf{z} \in \mathcal{L} + \mathbf{x}$. Now note that $\mathbf{w} \in F = F + \mathbf{0}$ and $\mathbf{w} = \mathbf{z} + (\mathbf{w} - \mathbf{z}) \in F + (\mathbf{w} - \mathbf{z})$. Therefore $(F + \mathbf{0}) \cap (F + \mathbf{w} - \mathbf{z}) \neq \emptyset$. But $\mathbf{0}$ and $\mathbf{w} - \mathbf{z}$ are distinct points in \mathcal{L} , a contradiction to our assumption on F . Hence $|(\mathcal{L} + \mathbf{x}) \cap F| \leq 1$ as needed. Assume F is \mathcal{L} -covering. Take $\mathbf{x} \in \mathbb{R}^n$. Since $\mathcal{L} + F = \mathbb{R}^n$, there exists $\mathbf{y} \in \mathcal{L}$ such that $\mathbf{x} \in \mathbf{y} + F$. Therefore $\mathbf{x} - \mathbf{y} \in F$, and since $-\mathbf{y} \in \mathcal{L}$ we

get that $|(\mathcal{L} + \mathbf{x}) \cap F| \geq 1$ as needed. The claim for F an \mathcal{L} -tiling follows directly from the previous assertions.

For a $A \subseteq \mathbb{R}^n$, we say that Λ is a tiling of A by F with respect to \mathcal{L} , if $A \subseteq \Lambda + F$, $\Lambda \subseteq \mathcal{L}$ and F is \mathcal{L} -tiling.

Lemma 2.4.3. *Let $B \in \mathbb{R}^{n \times n}$ denote a basis for a lattice \mathcal{L} . Then $F = B[0, 1]^n$ is a fundamental domain for \mathcal{L} and $\text{vol}_n(F) = \det(\mathcal{L})$.*

Proof. First, to compute the volume of F , we note that

$$\text{vol}_n(B[0, 1]^n) = |\det(B)| \text{vol}_n([0, 1]^n) = |\det(B)|.$$

Since $\det(\mathcal{L}) = \sqrt{\det(B^t B)} = \sqrt{\det(B)^2} = |\det(B)|$, we get the desired equality for volume F .

We prove that F is K -packing. Take distinct $\mathbf{x}, \mathbf{y} \in \mathcal{L}$. We wish to show that $\mathbf{x} + F \cap \mathbf{y} + F = \emptyset$. Now note that

$$\mathbf{x} + F \cap \mathbf{y} + F \neq \emptyset \Leftrightarrow \mathbf{x} - \mathbf{y} \notin F - F$$

Now applying B^{-1} to both sides, we get that $B^{-1}(\mathbf{x} - \mathbf{y}) \notin [0, 1]^n - [0, 1]^n = (-1, 1)^n$. Since $\mathbf{x} - \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$, we have that $B^{-1}(\mathbf{x} - \mathbf{y}) \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$. Clearly $\mathbb{Z}^n \setminus \{\mathbf{0}\} \cap (-1, 1)^n = \emptyset$ as needed.

We now show that F is \mathcal{L} -covering. Take $\mathbf{x} \in \mathbb{R}^n$. Let $\mathbf{z} = \lfloor B^{-1} \mathbf{x} \rfloor$ ($\lfloor \cdot \rfloor$ is performed component wise) and let $\mathbf{w} = B^{-1} \mathbf{x} - \mathbf{z}$. Then clearly $\mathbf{z} \in \mathbb{Z}^n$, and $\mathbf{w} \in [0, 1]^n$ (since $x - \lfloor x \rfloor \in [0, 1)$ for any $x \in \mathbb{R}$). But then

$$\mathbf{x} = B(\mathbf{z} + \mathbf{w}) = B\mathbf{z} + B\mathbf{w} \in \mathcal{L} + B[0, 1]^n = \mathcal{L} + F$$

as needed. □

Lemma 2.4.4. *Let $F \subseteq \mathbb{R}^n$ be measurable \mathcal{L} . Let $g : \mathbb{R}^n \rightarrow \mathbb{R}_+$ be a measurable function. If F is a \mathcal{L} -(packing, covering, tiling) we have that*

$$\int_F \sum_{\mathbf{y} \in \mathcal{L}} g(\mathbf{y} + \mathbf{x}) d\mathbf{x} \quad (\leq, \geq, =) \quad \int_{\mathbb{R}^n} g(\mathbf{x}) d\mathbf{x}.$$

Furthermore, if F is a \mathcal{L} -(packing, covering, tiling) we have that

$$\text{vol}_n(F) \quad (\leq, \geq, =) \quad \det(\mathcal{L}).$$

Proof. Since $g \geq 0$ and measurable, we have that the function $\mu(A) = \int_A g(\mathbf{x})d\mathbf{x}$, for $A \subseteq \mathbb{R}^n$ measurable, defines a measure on \mathbb{R}^n . Let $1_{\mathbf{y}+F}$, $\mathbf{y} \in \mathcal{L}$, denote the indicator function of $\mathbf{y} + F$. Since $\mathbf{y} + F$ is measurable, we get that $1_{F+\mathbf{y}}$ is non-negative measurable function. Since \mathcal{L} is countable, by the monotone convergence theorem we have that

$$\begin{aligned} \sum_{\mathbf{y} \in \mathcal{L}} \mu(\mathbf{y} + F) &= \sum_{\mathbf{y} \in \mathcal{L}} \int_{\mathbb{R}^n} 1_{\mathbf{y}+F}(\mathbf{x})g(\mathbf{x})d\mathbf{x} = \sum_{\mathbf{y} \in \mathcal{L}} \int_{\mathbb{R}^n} 1_F(\mathbf{x})g(\mathbf{x} + \mathbf{y})d\mathbf{x} \\ &= \int_{\mathbb{R}^n} \sum_{\mathbf{y} \in \mathcal{L}} 1_F(\mathbf{x})g(\mathbf{x} + \mathbf{y})d\mathbf{x} = \int_F \sum_{\mathbf{y} \in \mathcal{L}} g(\mathbf{x} + \mathbf{y})d\mathbf{x} \end{aligned}$$

If F is \mathcal{L} -packing, then collections of sets $\mathbf{y} + F \subseteq \mathbb{R}^n$, for $\mathbf{y} \in \mathcal{L}$, are all disjoint.

Therefore we have that

$$\int_{\mathbb{R}^n} g(\mathbf{x})d\mathbf{x} = \mu(\mathbb{R}^n) \geq \mu(\mathcal{L} + F) = \sum_{\mathbf{y} \in \mathcal{L}} \mu(\mathbf{y} + F) = \int_F \sum_{\mathbf{y} \in \mathcal{L}} g(\mathbf{x} + \mathbf{y})d\mathbf{x}$$

as needed. If F is \mathcal{L} -covering, we have that $\mathbb{R}^n \subseteq \mathcal{L} + F$, and hence

$$\mu(\mathbb{R}^n) = \mu(\mathcal{L} + F) \leq \sum_{\mathbf{y} \in \mathcal{L}} \mu(\mathbf{y} + F) = \int_F \sum_{\mathbf{y} \in \mathcal{L}} g(\mathbf{x} + \mathbf{y})d\mathbf{x}$$

as needed. If F is \mathcal{L} -tiling, we get the desired equality by combining the above two inequalities.

We now prove the furthermore. Let $B \in \mathbb{R}^{n \times n}$ denote a basis for \mathcal{L} . From Lemma 2.4.3, we know that $B[0, 1)^n$ is \mathcal{L} -tiling and satisfies $\text{vol}_n(B[0, 1)^n) = \det(\mathcal{L})$. From the first part of the lemma, we have that

$$\text{vol}_n(F) = \int_{\mathbb{R}^n} 1_F(\mathbf{x})d\mathbf{x} = \int_{B[0, 1)^n} \sum_{\mathbf{y} \in \mathcal{L}} 1_F(\mathbf{x} + \mathbf{y})d\mathbf{x} = \int_{B[0, 1)^n} |(\mathcal{L} + \mathbf{x}) \cap F|d\mathbf{x}$$

If F is \mathcal{L} -(packing, covering, tiling) we have that $\forall \mathbf{x} \in \mathbb{R}^n$, $|(\mathcal{L} + \mathbf{x}) \cap F|$ ($\leq, \geq, =$) 1.

Therefore if F is \mathcal{L} -(packing, covering, tiling) we have that

$$\text{vol}_n(F) = \int_{B[0, 1)^n} |(\mathcal{L} + \mathbf{x}) \cap F|d\mathbf{x} \quad (\leq, \geq, =) \quad \int_{B[0, 1)^n} 1d\mathbf{x} = \text{vol}_n(B[0, 1)^n) = \det(\mathcal{L})$$

as needed. \square

2.4.2 Lattice Inequalities

Perhaps the most fundamental inequality in the geometry of numbers is Minkowski's first theorem, which is stated as follows:

Theorem 2.4.5. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be an n dimensional lattice and let $K \subseteq \mathbb{R}^n$ denote a centrally symmetric convex body. Then*

$$\lambda_1(K, \mathcal{L}) \leq 2 \left(\frac{\det(\mathcal{L})}{\text{vol}(K)} \right)^{\frac{1}{n}}$$

Proof. Let $\lambda = \lambda_1(K, \mathcal{L})$. We claim that $\frac{\lambda}{2}\text{int}(K) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_K < \frac{\lambda}{2}\}$ is \mathcal{L} -packing. Take distinct $\mathbf{x}, \mathbf{y} \in \mathcal{L}$. Now note that

$$\begin{aligned} \mathbf{x} + \frac{\lambda}{2}\text{int}(K) \cap \mathbf{y} + \frac{\lambda}{2}\text{int}(K) \neq \emptyset &\Leftrightarrow \mathbf{x} - \mathbf{y} \in \frac{\lambda}{2}\text{int}(K) - \frac{\lambda}{2}\text{int}(K) \\ &\Leftrightarrow \mathbf{x} - \mathbf{y} \in \lambda\text{int}(K) \quad (\text{by symmetry of } K) \end{aligned}$$

Since $\mathbf{x} - \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$, we have that $\|\mathbf{x} - \mathbf{y}\|_K \geq \lambda$. Therefore $\mathbf{x} - \mathbf{y} \notin \lambda\text{int}(K)$ as needed. Since $\frac{\lambda}{2}\text{int}(K)$ is \mathcal{L} -packing, by Lemma 2.4.4 we have that

$$\text{vol}_n \left(\frac{\lambda}{2}\text{int}(K) \right) \leq \det(\mathcal{L}) \Leftrightarrow \lambda \leq \frac{2 \det(\mathcal{L})^{\frac{1}{n}}}{\text{vol}_n(\text{int}(K))^{\frac{1}{n}}}.$$

Since $\text{vol}_n(\text{int}(K)) = \text{vol}_n(K)$, the theorem follows. \square

The following lemma describes some classical relations between the different parameters of a lattice (see [25, 67]).

Lemma 2.4.6. *Let K denote an n dimensional convex body, and let \mathcal{L} denote an n -dimensional lattice. Then the following holds:*

- (1) $\lambda_n(K - K, \mathcal{L}) \leq \mu(K, \mathcal{L}) \leq \sum_{i=1}^n \lambda_i(K - K, \mathcal{L})$.
- (2) *If K is $\mathbf{0}$ -centered then $\lambda_{n-i}(K, \mathcal{L})\lambda_{i+1}(K^*, \mathcal{L}^*) \geq 1$, for $i \in [n]$.*

The following lemma strengthens the upper bound on $\mu(K, \mathcal{L})$ above due to Kannan and Lovasz [67].

Lemma 2.4.7. *Let $K \subseteq \mathbb{R}^n$ be a convex body, and \mathcal{L} be an n -dimensional lattice. Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ denote linearly independent \mathcal{L} , with gram schmidt vectors $\mathbf{b}_1^*, \dots, \mathbf{b}_n^*$. Then*

$$\mu(K, \mathcal{L}) \leq \sum_{i=1}^n \|\mathbf{b}_i^*\|_{\pi_i(K-K)}.$$

where π_i denotes the orthogonal projection map onto $\text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{i-1})$. Furthermore, if $\mathbf{b}_1, \dots, \mathbf{b}_n$ is a basis of \mathcal{L} then

$$\lambda_1(K - K, \mathcal{L}) \geq \min_{i \in [n]} \|\mathbf{b}_i^*\|_{\pi_i(K-K)}.$$

2.5 Computational Complexity

For a thorough treatment of the important concepts in algorithms and computational complexity, the interested reader may refer to the following references [123, 29].

In the study of algorithms, a first fundamental task is to understanding the running time of algorithms as the size of the input grows (i.e. asymptotic complexity). Other resources of interest will be the amount of space used by an algorithm, as well as amount of randomness used (i.e. number of random bits).

We introduce the big O notation for measuring asymptotic complexity. For functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$, we define the following asymptotic relations.

- (1) $f(n) = O(g(n))$ if $\exists C > 0, N_0 \in \mathbb{N}$ such that for $n \geq N_0$, $f(n) \leq Cg(n)$.
- (2) $f(n) = \Omega(g(n))$ if $\exists C > 0, N_0 \in \mathbb{N}$ such that for $n \geq N_0$, $f(n) \geq Cg(n)$.
- (3) $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f = \Omega(g(n))$.
- (4) $f(n) = o(g(n))$ if $\forall C > 0, \exists N_0 \in \mathbb{N}$ such that for $n \geq N_0$, $f(n) \leq Cg(n)$.

Instead of comparing the asymptotics of two functions, we will often want to compare one function to a class of functions. In computer science, one of the most important classes is the class $\text{poly}(n)$, i.e. the class of polynomials.

For $T : \mathbb{N} \rightarrow \mathbb{N}$ we say that $T(n) = \text{poly}(n)$ if $\exists k \in \mathbb{N}$ such that $T(n) = O(x^k)$. Let ALG be an algorithm which on an input of size n , runs in at most $T(n)$ time, and uses at most $S(n)$ space. We say that ALG runs in polynomial time if $T(n) = \text{poly}(n)$ and that ALG uses polynomial space if $S(n) = \text{poly}(n)$.

In this thesis, we will also be interested in algorithms that run in exponential time, where we denote this class of running times by the symbol $2^{O(n)}$. Here, we write $T(n) = 2^{O(n)}$, if there exists $c > 0$ such that $T(n) = O(2^{cn})$. As above, we will say that ALG runs in exponential time if $T(n) = 2^{O(n)}$, and uses exponential space if $S(n) = 2^{O(n)}$.

In computational complexity, one of the main algorithmic tasks is deciding membership in a language L . A language $L \subseteq \{0, 1\}^* = \bigcup_{n=1}^{\infty} \{0, 1\}^n$ is a collection (generally infinite) of $\{0, 1\}$ strings. We say that an algorithm ALG decides L if on input $\mathbf{x} \in \{0, 1\}^*$, $\text{ALG}(\mathbf{x})$ outputs 1 if $\mathbf{x} \in L$ and outputs 0 otherwise. ALG decides L in $T(n)$ time if for $\mathbf{x} \in \{0, 1\}^n$, $\text{ALG}(\mathbf{x})$ runs in time at most $T(n)$. We present the two most important classes of languages in computational complexity. We define the class P to be the set of languages which can be decided in polynomial time. More precisely, a language $L \in P$ if there exists an algorithm which decides L in polynomial time. We denote the class NP , to be the class of languages which can be verified in polynomial time. More precisely, $L \in NP$ if there exists a polynomial time verifier V satisfying the following:

- (1) If $\mathbf{x} \in L$, there exists $\mathbf{y} \in \{0, 1\}^*$ s.t. $V(\mathbf{x}, \mathbf{y}) = 1$.
- (2) If $\mathbf{x} \notin L$, then for all $\mathbf{y} \in \{0, 1\}^*$, $V(\mathbf{x}, \mathbf{y}) = 0$.

From the above definitions, it is easy to see that $P \subseteq NP$, however the class NP is thought to be significantly larger than P (this is the $P \neq NP$ problem). The canonical example of an NP language is the language of satisfiable boolean formulas (e.g. $x_1 \vee x_2 \vee x_3$), which we denote SAT . Starting with a boolean formula, given

an assignment to the variables one can easily check (i.e. in polynomial time) that the assignment satisfies the formula, and hence verify that the formula is in *SAT*. We therefore see that $SAT \in NP$ because the statement $\mathbf{x} \in SAT$ has a short and efficiently verifiable “proof” if it is true. The *SAT* problem is also important due to its relation to other problems in the class *NP*. In particular, *SAT* is understood to be one of the hardest problems in *NP*. The reason for this is that the membership problem for any language $L \in NP$ can be reduced in polynomial time to a question of membership in *SAT*. Hence we call *SAT* a complete problem for the class *NP*.

Many of the computational problems we examine in this thesis will in fact be either NP-Complete or NP-Hard (i.e. at least as hard as *SAT*), and hence some grounding in complexity theory is useful.

An important class of algorithms we will analyze in this thesis is the class of randomized algorithms. Randomized algorithms are algorithms which utilize random bits (i.e. random “coin flips”) during their executions to make decisions. An important feature of randomized algorithms is that they may not always succeed at the given computational task, i.e. they have a probability of failure. For example, when deciding a language L , given an input \mathbf{x} , an algorithm ALG may only correctly decide whether $\mathbf{x} \in L$ with 99% probability. We designate this class of randomized algorithms as the class of Monte Carlo algorithms. Another important class of randomized algorithms is the class of Las Vegas algorithms. A Las Vegas algorithm always guarantees that its output is correct, however the amount of computational resources used by the algorithm is a random variable. Hence for a Las Vegas algorithm, we generally only give bounds on the expected running time of the algorithm.

2.5.1 Computational Model

The algorithms presented in this thesis generally take in as part of their input a sequence of vectors or matrices with rational coefficients. Since complexity of the

algorithms will grow with the size of coefficients in these vectors and matrices, we will need to formally account for them. To do this, for a rational matrix $A \in \mathbb{Q}^{m \times n}$ we define $\text{enc}\langle A \rangle$ as the length of the binary encoding of A .

In this thesis, nearly all of our algorithms will need to interact in one way or another with a convex body or norm. To ensure our algorithms work in the most general settings, we will only require narrow types of oracle access to the body or norm in question. We define three different types of oracles that we will need for our algorithms. For convenience, norms here will always be indexed by their associated unit balls. With some slight modifications, we adopt the terminology from [56].

Let $K \subseteq \mathbb{R}^n$ be a convex body. For $\epsilon \geq 0$, we define

$$K^\epsilon = K + \epsilon B_2^n \quad \text{and} \quad K^{-\epsilon} = \{\mathbf{x} \in K : \mathbf{x} + \epsilon B_2^n \subseteq K\}$$

We say that K is (\mathbf{a}_0, R) -*circumscribed* if $K \subseteq \mathbf{a}_0 + RB_2^n$ for some $\mathbf{a}_0 \in \mathbb{Q}^n$ and $R \in \mathbb{Q}$. We say that K is (\mathbf{a}_0, r, R) -*centered* if $\mathbf{a}_0 + rB_2^n \subseteq K \subseteq \mathbf{a}_0 + RB_2^n$ for $\mathbf{a}_0 \in \mathbb{Q}^n$, $r, R \in \mathbb{Q}$. We will always assume that the above parameters are given explicitly as part of the input to our problems, and hence our algorithms will be allowed to depend on $\text{enc}\langle \mathbf{a}_0, r, R \rangle$.

Definition 2.5.1. A *weak membership oracle* O_K for K is function which takes as input a point $\mathbf{x} \in \mathbb{Q}^n$ and real $\epsilon > 0$, and returns

$$O_K(\mathbf{x}, \epsilon) = \begin{cases} 1 & : \mathbf{x} \in K^{-\epsilon} \\ 0 & : \mathbf{x} \notin K^\epsilon \end{cases}$$

where any answer is acceptable if $\mathbf{x} \in K^\epsilon \setminus K^{-\epsilon}$.

Definition 2.5.2. A *strong separation oracle* SEP_K for K on input $\mathbf{y} \in \mathbb{Q}^n$ either returns YES if $\mathbf{y} \in K$, or some $\mathbf{c} \in \mathbb{Q}^n$ such that $\langle \mathbf{c}, \mathbf{x} \rangle < \langle \mathbf{c}, \mathbf{y} \rangle$, $\forall \mathbf{x} \in K$.

When working with the above oracle, we assume that there is a polynomial ϕ , such that on input \mathbf{y} as above, the output of SEP_K has size bounded by $\phi(\text{enc}\langle \mathbf{y} \rangle)$. The runtimes of algorithms using SEP_K will therefore depend on ϕ .

Let K be a convex body containing the origin in its interior.

Definition 2.5.3. A *weak distance oracle* D_K for K is a function that takes as input a point $\mathbf{x} \in \mathbb{Q}^n$ and $\epsilon > 0$, and returns a rational number satisfying

$$|D_K(\mathbf{x}, \epsilon) - \|\mathbf{x}\|_K| \leq \epsilon \min\{1, \|\mathbf{x}\|_K\}.$$

As above, we assume the existence of a polynomial ϕ , such that the size of the output of D_K on (\mathbf{x}, ϵ) is bounded by $\phi(\text{enc}(\mathbf{x}, \epsilon))$. For a $(0, r, R)$ -centered body K , $\forall \mathbf{x} \in \mathbb{R}^n$, we crucially have that

$$\frac{1}{R}\|\mathbf{x}\| \leq \|\mathbf{x}\|_K \leq \frac{1}{r}\|\mathbf{x}\|.$$

In all the above oracles, we shall assume that $\dim(K) = n$, i.e. dimension of the ambient body, is encoded in unary in the guarantees of the oracles.

Definition 2.5.4. Let $f : K \rightarrow \mathbb{R}$ be a convex function with domain $K \subseteq \mathbb{R}^n$. f is equipped with a subgradient oracle, if we have query access to a subgradient $\mathbf{v} \in \partial f(\mathbf{x})$, for any $\mathbf{x} \in K$.

Oracle Time Complexity. To describe the running times and space requirements of our algorithms, we define the notion of oracle time and space complexity. We say that an algorithm runs in oracle $T(n)$ time and $S(n)$ space if it performs at most $T(n)$ arithmetic operations and calls to the oracle(s), and uses at most $S(n)$ space, on an input of size n . In our setting, we note that the input size includes the encoding length of all the oracle guarantees (i.e. inner / outer radius, etc). Since almost all our algorithms run in the oracle model, when we write that an algorithm runs in $T(n)$ time and $S(n)$ space we will always mean “oracle” $T(n)$ time and $S(n)$ space.

In this thesis, the complexity classes $T(n)$ and $S(n)$ will generally either denote the class $\text{poly}(n)$ or $2^{O(n)}$. Since many of the oracles used here return vector valued data (i.e. distance / separation oracles), we make the assumption that the output

complexity of the oracles is bounded by some polynomial ϕ (as mentioned above). Here we will guarantee that the claimed running times and space requirements are stable under any fixed choice of ϕ . Since our algorithms will take many parameters as input, for notational convenience we often describe their complexity by the expression $\text{poly}(\cdot)$ or $2^{O(n)} \text{poly}(\cdot)$ (here n will generally refer to the ambient dimension), where by $\text{poly}(\cdot)$ we mean polynomial in the length of all input parameters.

2.5.2 Operations on Convex Bodies

Here we summarize important results about the equivalence of certain oracles for a convex body, as well as how to build oracles for certain operations on a convex body. These results are explicitly given or easily derived from the results in [56].

Theorem 2.5.5. *Let $K \subseteq \mathbb{R}^n$ be a (\mathbf{a}_0, r, R) -centered convex body presented by a weak membership oracle O_K . The following oracles can be implemented from O_K in polynomial time:*

- (1) *A weak membership oracle for $(K - \mathbf{a}_0)^*$.*
- (2) *For $T \in \mathbb{Q}^{m \times n}$, $m \leq n$, T full rank, a weak membership oracle for TK .*

Theorem 2.5.6. *Let $K_1, K_2 \subseteq \mathbb{R}^n$ be centered convex bodies presented by weak membership oracles O_{K_1}, O_{K_2} . The following oracles can be implemented from O_{K_1} and O_{K_2} in polynomial time:*

- (1) *A weak membership oracle for $K_1 + K_2$.*
- (2) *A weak membership oracle for $\text{conv}\{K_1, K_2\}$.*
- (3) *Assuming $K_1 \cap K_2$ is centered, a weak membership oracle for $K_1 \cap K_2$.*

Lemma 2.5.7. *Let $K \subseteq \mathbb{R}^n$ be a $(0, r, R)$ -centered convex body. Then a weak distance oracle for $\|\cdot\|_K$ and a weak membership oracle for K are polynomial time equivalent.*

The following simple lemma allows us to construct a strong separation oracle for any hyperplane section of a convex body already equipped with a strong separation oracle.

Lemma 2.5.8. *Let $K \subseteq \mathbb{R}^n$ be a convex body presented by a strong separation oracle SEP_K . Let $H = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$ denote an affine subspace, where $A \in \mathbb{Q}^{m \times n}$, $\mathbf{b} \in \mathbb{Q}^m$. Then one can construct a separation oracle for $K \cap H$, such that on input $\mathbf{y} \in H$, the oracle executes in polynomial time using a single call to SEP_K .*

2.5.3 Fundamental Algorithms

Here we list some of the fundamental algorithmic tools we will require.

The following theorem yields the classical equivalence between weak membership and weak optimization [127, 56].

Theorem 2.5.9 (Convex Optimization via Ellipsoid Method). *Let $K \subseteq \mathbb{R}^n$ an (\mathbf{a}_0, r, R) -centered convex body given by a weak membership oracle O_K . Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ denote an L -Lipshitz convex function given by an oracle that, for every $\mathbf{x} \in \mathbb{Q}^n$ and $\delta > 0$, returns a rational number t such that $|f(\mathbf{x}) - t| \leq \delta$. Then for $\epsilon > 0$, a rational number ω and vector $\mathbf{y} \in K$ satisfying*

$$\omega - \epsilon \leq \min_{\mathbf{x} \in K} f(\mathbf{x}) \leq f(\mathbf{y}) \leq \omega$$

can be computed using O_K in polynomial time.

The following algorithm from [56], allows us to deterministically compute an ellipsoid with relatively good “sandwiching” guarantees for a convex body K .

Theorem 2.5.10 (Algorithm GLS-Round). *Let $K \subseteq \mathbb{R}^n$ be an (\mathbf{a}_0, R) -circumscribed convex body given by a strong-separation oracle SEP_K . Then for any $\epsilon > 0$, there is a polynomial time algorithm to compute $A \succ 0$, $A \in \mathbb{Q}^{n \times n}$ and $\mathbf{t} \in \mathbb{R}^n$, such that the ellipsoid $E = E(A)$ satisfies $K \subseteq E + \mathbf{t}$, and one of the following: (a) $\text{vol}_n(E) \leq \epsilon$, or (b) $\frac{1}{(n+1)n^{\frac{1}{2}}}E + \mathbf{t} \subseteq K$.*

The next algorithm comes from the literature on random walks on convex bodies [87, 86, 88]. The algorithm allows us to sample from essentially any logconcave measure on a convex body.

Theorem 2.5.11 (Algorithm Logconcave-Sampler, [86]). *Let $K \subseteq \mathbb{R}^n$ be a (\mathbf{a}_0, r, R) -centered convex body given by a weak membership oracle O_K . Let $f : K \rightarrow \mathbb{R}_+$ be a polynomial time computable log-concave function satisfying*

$$\sup_{\mathbf{x} \in K} f(\mathbf{x}) \leq \beta^n f(\mathbf{0})$$

for some $\beta > 1$. Let $\epsilon, \tau > 0$. Then the following can be computed:

(1) A random point $X \in K$ with distribution σ satisfying $d_{\text{TV}}(\sigma, \pi_f) \leq \tau$ polynomial time.

(2) A point $\mathbf{b} \in K$ and a matrix $A \in \mathbb{Q}^{n \times n}$ such that $\forall \mathbf{x} \in \mathbb{R}^n$

$$|\langle \mathbf{x}, \mathbf{b} - \mathbf{b}(f) \rangle| \leq \epsilon \mathbf{x}^t \text{cov}(f) \mathbf{x} \quad \text{and} \quad |\mathbf{x}^t (A - \text{cov}(f)) \mathbf{x}| \leq \epsilon \mathbf{x}^t \text{cov}(f) \mathbf{x},$$

with probability $1 - \delta$ in polynomial time.

CHAPTER III

ON THE CHVÁTAL-GOMORY CLOSURE OF A COMPACT CONVEX SET

In this Chapter, we show that the Chvátal-Gomory closure of any compact convex set is a rational polytope. This resolves an open question of Schrijver [119] for irrational polytopes¹, and generalizes the same result for the case of rational polytopes [119], rational ellipsoids [38] and strictly convex bodies [30]. This Chapter is based on the paper [34] (joint with Santanu Dey and Juan Pablo Vielma).

3.1 Introduction

Gomory [55] introduced the Gomory fractional cuts, also known as Chvátal-Gomory (CG) cuts [28], to design the first finite cutting plane algorithm for Integer Linear Programming (ILP). Since then, many important classes of facet-defining inequalities for combinatorial optimization problems have been identified as CG cuts. For example, the classical Blossom inequalities for general Matching [46] - which yield the integer hull - and Comb inequalities for the Traveling Salesman problem [57, 58] are both CG cuts over the base linear programming relaxations. CG cuts have also been effective from a computational perspective; see for example [19, 50]. Although CG cuts have traditionally been defined with respect to rational polyhedra for ILP, they straightforwardly generalize to the nonlinear setting and hence can also be used for convex Integer Nonlinear Programming (INLP), i.e. the class of discrete optimization problems whose continuous relaxation is a general convex optimization problem. CG

¹After the completion of this work, it has been brought to our notice that the polyhedrality of the Chvátal-Gomory Closure for irrational polytopes has recently been shown independently by J. Dunkel and A. S. Schulz in [43]. The proof presented in this Chapter has been obtained independently.

cuts for non-polyhedral sets were considered implicitly in [28, 119] and more explicitly in [26, 30, 38]. Let $K \subseteq \mathbb{R}^n$ be a closed convex set, and let h_K represent its support function. Given $\mathbf{a} \in \mathbb{Z}^n$, we define the CG cut for K derived from \mathbf{a} as the inequality

$$\langle \mathbf{a}, \mathbf{x} \rangle \leq \lfloor h_K(\mathbf{a}) \rfloor . \quad (3.1.1)$$

The CG closure of K is the convex set whose defining inequalities are exactly all the CG cuts for K . A classical result of Schrijver [119] is that the CG closure of a rational polyhedron is a rational polyhedron. Previously, we were able to verify that the CG closure of any strictly convex body² intersected with a rational polyhedron is a rational polyhedron [38, 30]. We remark that the proof requires techniques significantly different from those described in [119].

While the intersections of strictly convex bodies with rational polyhedra yield a large and interesting class of bodies, they do not capture many natural examples that arise in convex INLP. For example, it is not unusual for the feasible region of a semi-definite or conic-quadratic program [15] to have infinitely many faces of different dimensions, where additionally a majority of these faces cannot be isolated by intersecting the feasible region with a rational supporting hyperplane (as is the case for standard ILP with rational data). Roughly speaking, the main barrier to progress in the general setting has been a lack of understanding of how CG cuts act on irrational affine subspaces (affine subspaces whose defining equations cannot be described with rational data).

As a starting point for this study, perhaps the simplest class of bodies where current techniques break down are polytopes defined by irrational data. Schrijver considers these bodies in [119], and in a discussion section at the end of the paper, he writes ³:

²A full dimensional compact convex set whose only non-trivial faces are vertices. In this Chapter, we call zero dimensional faces as vertices.

³Theorem 1 in [119] is the result that the CG closure is a polyhedron. P' is the notation used for CG closure in [119]

“We do not know whether the analogue of Theorem 1 is true in real spaces. We were able to show only that if P is a bounded polyhedron in real space, and P' has empty intersection with the boundary of P , then P' is a (rational) polyhedron.”

In this Chapter, we prove that the CG closure of any compact convex set⁴ is a rational polytope, thus also resolving the question raised in [119]. As seen by Schrijver [119], most of the “action” in building the CG closure will indeed take place on the boundary of K . While the proof presented in this Chapter has some high level similarities to the one in [30], a substantially more careful approach was required to handle the general facial structure of a compact convex set (potentially infinitely many faces of all dimensions) and completely new ideas were needed to deal with faces having irrational affine hulls (including the whole body itself).

This Chapter is organized as follows. In Section 3.2 we introduce some notation, formally state our main result and give an overview of the proof. We then proceed with the full proof which is presented in Sections 3.3–3.5.

3.2 Definitions, Main Result and Proof Idea

Definition 3.2.1 (CG Closure). For a convex set $K \subseteq \mathbb{R}^n$ and $S \subseteq \mathbb{Z}^n$ let $CG(K, S) := \bigcap_{\mathbf{y} \in S} \{x \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle \leq \lfloor h_K(\mathbf{y}) \rfloor\}$. The CG closure of K is defined to be the set $CG(K) := CG(K, \mathbb{Z}^n)$.

The following theorem is the main result of the Chapter.

Theorem 3.2.2. *If $K \subseteq \mathbb{R}^n$ is a non-empty compact convex set, then $CG(K)$ is finitely generated. That is, there exists $S \subseteq \mathbb{Z}^n$ such that $|S| < \infty$ and $CG(K) = CG(K, S)$. In particular $CG(K)$ is a rational polyhedron.*

⁴If the convex hull of integer points in a convex set is not polyhedral, then the CG closure cannot be expected to be polyhedral. Since we do not have a good understanding of when this holds for unbounded convex sets, we restrict our attention here to the CG closure of compact convex sets.

For basic definitions related to convexity (i.e. faces, halfspaces, support function, etc.), we refer the reader to sections 2.1.1 and 2.2. We will use the following additional definitions and notation: For a convex set K and $\mathbf{v} \in \mathbb{R}^n$, let $H_{\mathbf{v}}^{\leq}(K) := \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{v}, \mathbf{x} \rangle \leq h_K(\mathbf{v})\}$ denote the supporting halfspace defined by \mathbf{v} for K , and let $H_{\mathbf{v}}(K) := \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{v}, \mathbf{x} \rangle = h_K(\mathbf{v})\}$ denote the supporting hyperplane. Let $F_{\mathbf{v}}(K) := K \cap H_{\mathbf{v}}(K)$ denote the face of K exposed by \mathbf{v} . If the context is clear, then we drop the K and simply write $H_{\mathbf{v}}^{\leq}$, $H_{\mathbf{v}}$ and $F_{\mathbf{v}}$. For $A \subseteq \mathbb{R}^n$, let $\text{aff}(A)$ denote the smallest affine subspace containing A . Furthermore let $\text{aff}_I(A) := \text{aff}(\text{aff}(A) \cap \mathbb{Z}^n)$, i.e. the largest integer subspace in $\text{aff}(A)$.

We present the outline of the proof for Theorem 3.2.2. The proof proceeds by induction on the dimension of K . The base case (K is a single point) is trivial. By the induction hypothesis, we can assume that (\dagger) every proper exposed face of K has a finitely generated CG closure. We build the CG closure of K in stages, proceeding as follows:

- (1) (Section 3.3) For $F_{\mathbf{v}}$, a proper exposed face, where $\mathbf{v} \in \mathbb{R}^n$, show that $\exists S \subseteq \mathbb{Z}^n$, $|S| < \infty$ such that $CG(K, S) \cap H_{\mathbf{v}} = CG(F_{\mathbf{v}})$ and $CG(K, S) \subseteq H_{\mathbf{v}}^{\leq}$ using (\dagger) and by proving the following:
 - (a) (Section 3.3.1) A CG cut for $F_{\mathbf{v}}$ can be rotated or “lifted” to a CG cut for K such that points in $F_{\mathbf{v}} \cap \text{aff}_I(H_{\mathbf{v}})$ separated by the original CG cut for $F_{\mathbf{v}}$ are separated by the new “lifted” one.
 - (b) (Section 3.3.2) A finite number of CG cuts for K separate all points in $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$ and all points in $\mathbb{R}^n \setminus H_{\mathbf{v}}^{\leq}$.
- (2) (Section 3.4) Create an approximation $CG(K, S)$ of $CG(K)$ such that (i) $|S| < \infty$, (ii) $CG(K, S) \subseteq K \cap \text{aff}_I(K)$ (iii) $CG(K, S) \cap \text{relbd}(K) = CG(K) \cap \text{relbd}(K)$. This is done in two steps:

- (a) (Section 3.4.1) Using the lifted CG closures of $F_{\mathbf{v}}$ from (1.) and a compactness argument on the sphere, create a first approximation $CG(K, S)$ satisfying (i) and (ii).
 - (b) (Section 3.4.2) Noting that $CG(K, S) \cap \text{relbd}(K)$ is contained in the union of a finite number of proper exposed faces of K , add the lifted CG closures for each such face to S to satisfy (iii).
- (3) (Section 3.5) We establish the final result by showing that there are only a finite number of CG cuts which separate at least one vertex of the approximation of the CG closure from (2).

3.3 $CG(K, S) \cap H_{\mathbf{v}} = CG(F_{\mathbf{v}})$ **and** $CG(K, S) \subseteq H_{\mathbf{v}}^{\leq}$

When K is a rational polyhedron, a key property of the CG closure is that for every face F of K , we have that (*) $CG(F) = F \cap CG(K)$. In this setting, a relatively straightforward induction argument coupled with (*) allows one to construct the approximation of the CG closure described above. In our setting, where K is compact convex, the approach taken is similar in spirit, though we will encounter significant difficulties. First, since K can have infinitely many faces, we must couple our induction with a careful compactness argument. Second and more significantly, establishing (*) for compact convex sets is substantially more involved than for rational polyhedra. As we will see in the following sections, the standard lifting argument to prove (*) for rational polyhedra cannot be used directly and must be replaced by a more involved two stage argument.

3.3.1 Lifting CG Cuts

To prove $CG(F) = F \cap CG(K)$ one generally uses a ‘lifting approach’, i.e., given a CG cut $CG(F, \{\mathbf{w}\})$ for F , $\mathbf{w} \in \mathbb{Z}^n$, we show that there exists a CG cut $CG(K, \{\mathbf{w}'\})$

for K , $\mathbf{w}' \in \mathbb{Z}^n$, such that

$$CG(K, \{\mathbf{w}'\}) \cap \text{aff}(F) \subseteq CG(F, \{\mathbf{w}\}) \cap \text{aff}(F). \quad (3.3.1)$$

To prove (3.3.1) when K is a rational polyhedron, one proceeds as follows. For the face F of K , we compute $\mathbf{v} \in \mathbb{Z}^n$ such that $F_{\mathbf{v}}(K) = F$ and $h_K(\mathbf{v}) \in \mathbb{Z}$. For $\mathbf{w} \in \mathbb{Z}^n$, we return the lifting $\mathbf{w}' = \mathbf{w} + l\mathbf{v}$, $l \in \mathbb{Z}_{>0}$, where l is chosen such that $h_K(\mathbf{w}') = h_F(\mathbf{w}')$. For general convex bodies though, neither of these steps may be achievable. When K is strictly convex however, in [30] we show that the above procedure can be generalized. First, every proper face F of K is an exposed vertex, hence $\exists \mathbf{x} \in K, \mathbf{v} \in \mathbb{R}^n$ such that $F = F_{\mathbf{v}} = \{\mathbf{x}\}$. For $\mathbf{w} \in \mathbb{Z}^n$, we show that setting $\mathbf{w}' = \mathbf{w} + \mathbf{v}'$, where \mathbf{v}' is a fine enough Dirichlet approximation (see Theorem 3.3.4 below) to a scaling of \mathbf{v} is sufficient for (3.3.1). In the proof, we critically use that F is simply a vertex. In the general setting, when K is a compact convex set, we can still meaningfully lift CG cuts, but not from all faces and not with exact containment. First, we only guarantee lifting for an exposed face $F_{\mathbf{v}}$ of K . Second, when lifting a CG cut for $F_{\mathbf{v}}$ derived from $\mathbf{w} \in \mathbb{Z}^n$, we only guarantee the containment on $\text{aff}_I(H_{\mathbf{v}})$, i.e. $CG(K, \mathbf{w}') \cap \text{aff}_I(H_{\mathbf{v}}) \subseteq CG(F, \mathbf{w}) \cap \text{aff}_I(H_{\mathbf{v}})$. This lifting, Proposition 3.3.5 below, uses the same Dirichlet approximation technique as in [30] but with a more careful analysis. Since we only guarantee the behavior of the lifting \mathbf{w}' on $\text{aff}_I(H_{\mathbf{v}})$, we will have to deal with the points in $\text{aff}(F) \setminus \text{aff}_I(H_{\mathbf{v}})$ separately, which we discuss in the next section.

Lemmas 3.3.1- 3.3.3 are technical results that are needed for proving Proposition 3.3.5.

Lemma 3.3.1. *Let K be a compact convex set in \mathbb{R}^n . Let $\mathbf{v} \in \mathbb{R}^n$, and let $(\mathbf{x}_i)_{i=1}^{\infty}$, $\mathbf{x}_i \in K$, be a sequence such that $\lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_i \rangle = h_K(\mathbf{v})$. Then*

$$\lim_{i \rightarrow \infty} d(F_{\mathbf{v}}(K), \mathbf{x}_i) = 0.$$

Proof. Let us assume that $\lim_{i \rightarrow \infty} d(F_{\mathbf{v}}(K), \mathbf{x}_i) \neq 0$. Then there exists an $\epsilon > 0$ such that for some subsequence $(\mathbf{x}_{\alpha_i})_{i=1}^{\infty}$ of $(\mathbf{x}_i)_{i=1}^{\infty}$ we have that $d(F_{\mathbf{v}}(K), \mathbf{x}_{\alpha_i}) \geq \epsilon$. Since $(\mathbf{x}_{\alpha_i})_{i=1}^{\infty}$ is an infinite sequence on a compact set K , there exists a convergent subsequence $(\mathbf{x}_{\beta_i})_{i=1}^{\infty}$ where $\lim_{i \rightarrow \infty} \mathbf{x}_{\beta_i} = \mathbf{x}$ and $\mathbf{x} \in K$. Now we note that $d(F_{\mathbf{v}}(K), \mathbf{x}) = \lim_{i \rightarrow \infty} d(F_{\mathbf{v}}(K), \mathbf{x}_{\beta_i}) \geq \epsilon$, where the first equality follows from the continuity of $d(F_{\mathbf{v}}(K), \cdot)$. Since $d(F_{\mathbf{v}}(K), \mathbf{x}) > 0$ we have that $\mathbf{x} \notin F_{\mathbf{v}}(K)$. On the other hand,

$$h_K(\mathbf{v}) = \lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_i \rangle = \lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_{\beta_i} \rangle = \langle \mathbf{v}, \mathbf{x} \rangle$$

and hence $\mathbf{x} \in F_{\mathbf{v}}(K)$, a contradiction. \square

Lemma 3.3.2. *Let K be a compact convex set in \mathbb{R}^n . Let $\mathbf{v} \in \mathbb{R}^n$, and let $(\mathbf{v}_i)_{i=1}^{\infty}$, $\mathbf{v}_i \in \mathbb{R}^n$, be a sequence such that $\lim_{i \rightarrow \infty} \mathbf{v}_i = \mathbf{v}$. Then for any sequence $(\mathbf{x}_i)_{i=1}^{\infty}$, $\mathbf{x}_i \in F_{\mathbf{v}_i}(K)$, we have that*

$$\lim_{i \rightarrow \infty} d(F_{\mathbf{v}_i}(K), \mathbf{x}_i) = 0.$$

Proof. We claim that $\lim_{i \rightarrow \infty} \langle \mathbf{x}_i, \mathbf{v} \rangle = h_K(\mathbf{v})$. Since K is compact, there exists $R \geq 0$ such that $K \subseteq RB^n$. Hence we get that

$$\begin{aligned} h_K(\mathbf{v}) &= \lim_{i \rightarrow \infty} h_K(\mathbf{v}_i) = \lim_{i \rightarrow \infty} \langle \mathbf{v}_i, \mathbf{x}_i \rangle \\ &= \lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_i \rangle + \langle \mathbf{v}_i - \mathbf{v}, \mathbf{x}_i \rangle \leq \lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_i \rangle + \|\mathbf{v}_i - \mathbf{v}\|R = \lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_i \rangle, \end{aligned}$$

where the first equality follows by continuity of h_K (h_K is convex on \mathbb{R}^n and finite valued). Since each $\mathbf{x}_i \in K$, we get the opposite inequality $\lim_{i \rightarrow \infty} \langle \mathbf{v}, \mathbf{x}_i \rangle \leq h_K(\mathbf{v})$ and hence we get equality throughout. Now by Lemma 3.3.1 we get that $\lim_{i \rightarrow \infty} d(F_{\mathbf{v}_i}(K), \mathbf{x}_i) = 0$ as needed. \square

The next lemma describes the central mechanics of the lifting process explained above. The sequence $(\mathbf{w}_i)_{i=1}^{\infty}$ will eventually denote the sequence of Dirichlet approximates of the scaling of \mathbf{v} added to \mathbf{w} , where one of these will serve as the lifting

w' .

Lemma 3.3.3. *Let $K \subseteq \mathbb{R}^n$ be a compact convex set. Take $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, $\mathbf{v} \neq 0$. Let $(\mathbf{w}_i, t_i)_{i=1}^\infty$, $\mathbf{w}_i \in \mathbb{R}^n$, $t_i \in \mathbb{R}_+$ be a sequence such that*

$$a. \lim_{i \rightarrow \infty} t_i = \infty, \quad b. \lim_{i \rightarrow \infty} \mathbf{w}_i - t_i \mathbf{v} = \mathbf{w}. \quad (3.3.2)$$

Then for every $\epsilon > 0$ there exists $N_\epsilon \geq 0$ such that for all $i \geq N_\epsilon$

$$h_K(\mathbf{w}_i) + \epsilon \geq t_i h_K(\mathbf{v}) + h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \geq h_K(\mathbf{w}_i) - \epsilon. \quad (3.3.3)$$

Proof. By (3.3.2) a,b we have that

$$\lim_{i \rightarrow \infty} \frac{\mathbf{w}_i}{t_i} = \mathbf{v} \quad (3.3.4)$$

and that we may pick $N_1 \geq 0$ such that

$$\|\mathbf{w}_i - t_i \mathbf{v}\| \leq \|\mathbf{w}\| + 1 \leq C \quad \text{for } i \geq N_1. \quad (3.3.5)$$

Let $(\mathbf{x}_i)_{i=1}^\infty$ be any sequence such that $\mathbf{x}_i \in F_{\mathbf{w}_i}(K) = F_{\mathbf{w}_i/t_i}(K)$. For each $i \geq 1$, let $\tilde{\mathbf{x}}_i = \arg \min_{\mathbf{y} \in F_{\mathbf{v}}(K)} \|\mathbf{x}_i - \mathbf{y}\|$. By (3.3.4) and Lemma 3.3.2, we may pick $N_2 \geq 0$ such that

$$d(F_{\mathbf{v}}(K), \mathbf{x}_i) = \|\mathbf{x}_i - \tilde{\mathbf{x}}_i\| \leq \frac{\epsilon}{2C} \quad \text{for } i \geq N_2. \quad (3.3.6)$$

Since $h_{F_{\mathbf{v}}(K)}$ is a continuous function, we may pick $N_3 \geq 0$ such that

$$|h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i - t_i \mathbf{v}) - h_{F_{\mathbf{v}}(K)}(\mathbf{w})| \leq \frac{\epsilon}{2} \quad \text{for } i \geq N_3. \quad (3.3.7)$$

Let $N_\epsilon = \max\{N_1, N_2, N_3\}$. Now since $\mathbf{x}_i \in F_{\mathbf{w}_i}(K)$ and $\tilde{\mathbf{x}}_i \in F_{\mathbf{v}}(K)$ we have that

$$\langle \mathbf{x}_i, \mathbf{w}_i \rangle \geq \langle \tilde{\mathbf{x}}_i, \mathbf{w}_i \rangle \quad \text{and} \quad \langle \tilde{\mathbf{x}}_i, t_i \mathbf{v} \rangle \geq \langle \mathbf{x}_i, t_i \mathbf{v} \rangle. \quad (3.3.8)$$

From (3.3.5), (3.3.6), (3.3.8) we get that for $i \geq N_\epsilon$

$$\begin{aligned} \langle \mathbf{x}_i, \mathbf{w}_i \rangle - \langle \tilde{\mathbf{x}}_i, \mathbf{w}_i \rangle &\leq \langle \mathbf{x}_i, \mathbf{w}_i \rangle - \langle \tilde{\mathbf{x}}_i, \mathbf{w}_i \rangle + \langle \tilde{\mathbf{x}}_i, t_i \mathbf{v} \rangle - \langle \mathbf{x}_i, t_i \mathbf{v} \rangle = \langle \mathbf{x}_i - \tilde{\mathbf{x}}_i, \mathbf{w}_i - t_i \mathbf{v} \rangle \\ &\leq \|\mathbf{x}_i - \tilde{\mathbf{x}}_i\| \|\mathbf{w}_i - t_i \mathbf{v}\| \leq \left(\frac{\epsilon}{2C}\right) C = \frac{\epsilon}{2}. \end{aligned} \quad (3.3.9)$$

From (3.3.9) we see that for $i \geq N_\epsilon$

$$h_K(\mathbf{w}_i) \geq h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i) \geq \langle \mathbf{w}_i, \tilde{\mathbf{x}}_i \rangle \geq \langle \mathbf{w}_i, \mathbf{x}_i \rangle - \frac{\epsilon}{2} = h_K(\mathbf{w}_i) - \frac{\epsilon}{2}. \quad (3.3.10)$$

Since $\langle \mathbf{v}, \cdot \rangle$ is constant on $F_{\mathbf{v}}(K)$, we have that

$$\begin{aligned} h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i) &= h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i - t_i \mathbf{v} + t_i \mathbf{v}) = h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i - t_i \mathbf{v}) + t_i h_{F_{\mathbf{v}}(K)}(\mathbf{v}) \\ &= h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i - t_i \mathbf{v}) + t_i h_K(\mathbf{v}) \end{aligned} \quad (3.3.11)$$

Combining (3.3.7), (3.3.10) and (3.3.11) we get that for $i \geq N_\epsilon$,

$$h_K(\mathbf{w}_i) + \epsilon \geq t_i h_K(\mathbf{v}) + h_{F_{\mathbf{v}}(K)}(\mathbf{w}_i) \geq h_K(\mathbf{w}_i) - \epsilon$$

as needed. □

Theorem 3.3.4 (Dirichlet's Approximation Theorem). *Let $(\alpha_1, \dots, \alpha_l) \in \mathbb{R}^l$. Then for every positive integer N , there exists $1 \leq n \leq N$ such that $\max_{1 \leq i \leq l} |n\alpha_i - \lfloor n\alpha_i \rfloor| \leq 1/N^{1/l}$.*

Proposition 3.3.5. *Let $K \subseteq \mathbb{R}^n$ be a compact and convex set, $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{w} \in \mathbb{Z}^n$. Then $\exists \mathbf{w}' \in \mathbb{Z}^n$ such that $CG(K, \mathbf{w}') \cap \text{aff}_I(H_{\mathbf{v}}(K)) \subseteq CG(F_{\mathbf{v}}(K), \mathbf{w}) \cap \text{aff}_I(H_{\mathbf{v}}(K))$.*

Proof. First, by possibly multiplying \mathbf{v} by a positive scalar we may assume that $h_K(\mathbf{v}) \in \mathbb{Z}$. Let $S = \text{aff}_I(H_{\mathbf{v}}(K))$. We may assume that $S \neq \emptyset$, since otherwise the statement is trivially true.

From Theorem 3.3.4 for any $\mathbf{v} \in \mathbb{R}^n$ there exists $(\mathbf{s}_i, t_i)_{i=1}^\infty$, $\mathbf{s}_i \in \mathbb{Z}^n$, $t_i \in \mathbb{N}$ such that (a.) $t_i \rightarrow \infty$ and (b.) $\|\mathbf{s}_i - t_i \mathbf{v}\| \rightarrow 0$. Now define the sequence $(\mathbf{w}_i, t_i)_{i=1}^\infty$, where $\mathbf{w}_i = \mathbf{w} + \mathbf{s}_i$, $i \geq 1$. Note that the sequence (\mathbf{w}_i, t_i) satisfies (3.3.2) and hence by Lemma 3.3.3 for any $\epsilon > 0$, there exists N_ϵ such that (3.3.3) holds. Let $\epsilon = \frac{1}{2}(1 - (h_{F_{\mathbf{v}}(K)}(\mathbf{w}) - \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \rfloor))$, and let $N_1 = N_\epsilon$. Note that $\lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) + \epsilon \rfloor = \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \rfloor$. Hence, since $h_K(\mathbf{v}) \in \mathbb{Z}$ by assumption, for all $i \geq N_1$ we have that

$$\lfloor h_K(\mathbf{w}_i) \rfloor \leq \lfloor t_i h_K(\mathbf{v}) + h_{F_{\mathbf{v}}(K)}(\mathbf{w}) + \epsilon \rfloor = t_i h_K(\mathbf{v}) + \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) + \epsilon \rfloor = t_i h_K(\mathbf{v}) + \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \rfloor.$$

Now pick $\mathbf{z}_1, \dots, \mathbf{z}_k \in S \cap \mathbb{Z}^n$ such that $\text{aff}(\mathbf{z}_1, \dots, \mathbf{z}_k) = S$ and let $R = \max\{\|\mathbf{z}_j\| : 1 \leq j \leq k\}$. Choose N_2 such that $\|\mathbf{w}_i - t_i \mathbf{v} - \mathbf{w}\| \leq \frac{1}{2R}$ for $i \geq N_2$. Now note that for $i \geq N_2$, $|\langle \mathbf{z}_j, \mathbf{w}_i \rangle - \langle \mathbf{z}_j, t_i \mathbf{v} + \mathbf{w} \rangle| = |\langle \mathbf{z}_j, \mathbf{w}_i - t_i \mathbf{v} - \mathbf{w} \rangle| \leq \|\mathbf{z}_j\| \|\mathbf{w}_i - t_i \mathbf{v} - \mathbf{w}\| \leq R \frac{1}{2R} = \frac{1}{2} \quad \forall j \in \{1, \dots, k\}$.

Next note that since $\mathbf{z}_j, \mathbf{w}_i \in \mathbb{Z}^n$, $\langle \mathbf{z}_j, \mathbf{w}_i \rangle \in \mathbb{Z}$. Furthermore, $t_i \in \mathbb{N}$, $\langle \mathbf{v}, \mathbf{z}_j \rangle = h_K(\mathbf{v}) \in \mathbb{Z}$ and $\mathbf{w} \in \mathbb{Z}^n$ implies that $\langle \mathbf{z}_j, t_i \mathbf{v} + \mathbf{w} \rangle \in \mathbb{Z}$. Given this, we must have $\langle \mathbf{z}_j, \mathbf{w}_i \rangle = \langle \mathbf{z}_j, t_i \mathbf{v} + \mathbf{w} \rangle \quad \forall j \in \{1, \dots, k\}, i \geq 1$ and hence we get $\langle \mathbf{x}, \mathbf{w}_i \rangle = \langle \mathbf{x}, t_i \mathbf{v} + \mathbf{w} \rangle \quad \forall \mathbf{x} \in S, i \geq 1$.

Let $\mathbf{w}' = \mathbf{w}_i$ where $i = \max\{N_1, N_2\}$. Now examine the set

$$L = \{\mathbf{x} : \langle \mathbf{x}, \mathbf{w}' \rangle \leq \lfloor h_K(\mathbf{w}') \rfloor\} \cap S.$$

Here we get that $\langle \mathbf{x}, \mathbf{w}_i \rangle \leq t_i h_K(\mathbf{v}) + \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \rfloor$ and $\langle \mathbf{x}, \mathbf{v} \rangle = h_K(\mathbf{v})$ for all $\mathbf{x} \in L$. Hence, we see that $\langle \mathbf{x}, \mathbf{w}_i - t_i \mathbf{v} \rangle \leq \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \rfloor$ for all $\mathbf{x} \in L$. Furthermore, since $\langle \mathbf{x}, \mathbf{w}_i - t_i \mathbf{v} \rangle = \langle \mathbf{x}, \mathbf{w} \rangle$ for all $\mathbf{x} \in L \subseteq S$, we have that $\langle \mathbf{x}, \mathbf{w} \rangle \leq \lfloor h_{F_{\mathbf{v}}(K)}(\mathbf{w}) \rfloor$ for all $\mathbf{x} \in L$, as needed. \square

3.3.2 Separating All Points in $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$

Since the guarantees on the lifted CG cuts produced in the previous section are restricted to $\text{aff}_I(H_{\mathbf{v}})$, we must still deal with the points in $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$. In this section, we show that points in $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$ can be separated by using a finite number of CG cuts in Proposition 3.3.9. To prove this, we will need Kronecker's theorem on simultaneous diophantine approximation which is stated next. See Niven [100] or Cassels [24] for a proof.

Theorem 3.3.6. *Let $(x_1, \dots, x_n) \in \mathbb{R}^n$ be such that the numbers $x_1, \dots, x_n, 1$ are linearly independent over \mathbb{Q} . Then the set $\{(nx_1 \pmod{1}, \dots, nx_n \pmod{1}) : n \in \mathbb{N}\}$ is dense in $[0, 1)^n$.*

The following lemmas allow us to normalize the vector \mathbf{v} defining $F_{\mathbf{v}}$ and $H_{\mathbf{v}}$ and simplify the analysis that follows.

Lemma 3.3.7. *Let $K \subseteq \mathbb{R}^n$ be a closed convex set, and let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be an invertible linear transformation. Then $h_K(\mathbf{v}) = h_{TK}(T^{-t}\mathbf{v})$ and $F_{\mathbf{v}}(K) = T^{-1}(F_{T^{-t}\mathbf{v}}(TK))$ for all $\mathbf{v} \in \mathbb{R}^n$. Furthermore, if T is a unimodular transformation, then $CC(K) = T^{-1}(CC(TK))$.*

Proof. Observe that

$$h_{TK}(T^{-t}\mathbf{v}) = \sup_{\mathbf{x} \in TK} \langle T^{-t}\mathbf{v}, \mathbf{x} \rangle = \sup_{\mathbf{x} \in K} \langle T^{-t}\mathbf{v}, T\mathbf{x} \rangle = \sup_{\mathbf{x} \in K} \langle \mathbf{v}, \mathbf{x} \rangle = h_K(\mathbf{v}).$$

Now note that

$$\begin{aligned} T^{-1}(F_{T^{-t}\mathbf{v}}(TK)) &= T^{-1}(\{\mathbf{x} : \mathbf{x} \in TK, h_{TK}(T^{-t}\mathbf{v}) = \langle T^{-t}\mathbf{v}, \mathbf{x} \rangle\}) \\ &= \{\mathbf{x} : T\mathbf{x} \in TK, h_{TK}(T^{-t}\mathbf{v}) = \langle T^{-t}\mathbf{v}, T\mathbf{x} \rangle\} \\ &= \{\mathbf{x} : \mathbf{x} \in K, h_K(\mathbf{v}) = \langle \mathbf{v}, \mathbf{x} \rangle\} = F_{\mathbf{v}}(K). \end{aligned}$$

Finally,

$$\begin{aligned} T^{-1}(CC(TK)) &= T^{-1}(\{\mathbf{x} : \mathbf{x} \in TK, \langle \mathbf{v}, \mathbf{x} \rangle \leq \lfloor h_{TK}(\mathbf{v}) \rfloor \forall \mathbf{v} \in \mathbb{Z}^n\}) \\ &= \{\mathbf{x} : T\mathbf{x} \in TK, \langle \mathbf{v}, T\mathbf{x} \rangle \leq \lfloor h_{TK}(\mathbf{v}) \rfloor \forall \mathbf{v} \in \mathbb{Z}^n\} \\ &= \{\mathbf{x} : T\mathbf{x} \in TK, \langle T^{-t}\mathbf{v}, T\mathbf{x} \rangle \leq \lfloor h_{TK}(T^{-t}\mathbf{v}) \rfloor \forall \mathbf{v} \in \mathbb{Z}^n\} \\ &= \{\mathbf{x} : \mathbf{x} \in K, \langle \mathbf{v}, \mathbf{x} \rangle \leq \lfloor h_K(\mathbf{v}) \rfloor \forall \mathbf{v} \in \mathbb{Z}^n\} = CC(K). \end{aligned}$$

□

Lemma 3.3.8. *Take $\mathbf{v} \in \mathbb{R}^n$. Then there exists an unimodular transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\lambda \in \mathbb{Q}_{>0}$ such that for $\mathbf{v}' = \lambda T\mathbf{v}$ we get that*

$$\mathbf{v}' = \left(\underbrace{0, \dots, 0}_t, \underbrace{1}_s, \alpha_1, \dots, \alpha_r \right), \quad (3.3.12)$$

where $t, r \in \mathbb{Z}_+$, $s \in \{0, 1\}$, and $\{1, \alpha_1, \dots, \alpha_r\}$ are linearly independent over \mathbb{Q} . Furthermore, we have that $\mathcal{D}(\mathbf{v}) = \inf\{\dim(W) : \mathbf{v} \in W, W = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = 0\}, A \in \mathbb{Q}^{m \times n}\} = s + r$.

Proof. Choose a permutation matrix P such that the rational entries of Pa form a contiguous block starting from the first entry of $P\mathbf{a}$, i.e. let $k \in \{0, \dots, n\}$ such that $(P\mathbf{a})_1, \dots, (P\mathbf{a})_k \in \mathbb{Q}$ and $(P\mathbf{a})_{k+1}, \dots, (P\mathbf{a})_n \in \mathbb{R} \setminus \mathbb{Q}$. Now we set our initial transformation $T \leftarrow P$, $\lambda \leftarrow 1$, and working vector $\mathbf{a}' \leftarrow P\mathbf{a}$. In what follows, we will apply successive updates to T, λ and \mathbf{a}' such that we maintain that T is unimodular, $\lambda \in \mathbb{Q}_{>0}$, and $\mathbf{a}' = \lambda T\mathbf{a}$.

First consider a vector $\mathbf{a}' \in \mathbb{R}^n$ such that a'_1, \dots, a'_k are rational and $(1, a'_{k+1}, \dots, a'_n)$ are linearly independent over \mathbb{Q} . If $k = 0$, i.e. $(1, a'_1, \dots, a'_n)$ are linearly independent over \mathbb{Q} , then we are done. We may therefore assume that $k \geq 1$. Similarly, if $(a'_1, \dots, a'_k) = 0^k$, then again we are done. Now let $\mathbf{a}'_R = (a'_1, \dots, a'_k)$ and $\mathbf{a}'_I = (a'_{k+1}, \dots, a'_n)$. By our assumptions, we note that $\mathbf{a}'_R \neq 0$. Via an appropriate scaling $\lambda' \in \mathbb{Q}_{>0}$, we may achieve $\lambda'\mathbf{a}'_R \in \mathbb{Z}^k$ and $\gcd(\lambda'a'_1, \dots, \lambda'a'_k) = 1$. Since $\lambda' \in \mathbb{Q}$, note that $(1, a'_{k+1}, \dots, a'_n)$ are linearly independent over \mathbb{Q} iff $(1, \lambda'a'_{k+1}, \dots, \lambda'a'_n)$ are. Set $\lambda \leftarrow \lambda'\lambda$ and $\mathbf{a}' \leftarrow \lambda'\mathbf{a}'$. Next, applying the Euclidean algorithm on the vector \mathbf{a}'_R , we get a unimodular transformation E such that

$$E\mathbf{a}'_R = (0^{k-1}, \gcd(a'_1, \dots, a'_k)) = (0^{k-1}, 1).$$

Now define the unimodular transformation T' , where

$$T'(\mathbf{x}) = (E(\mathbf{x}_1, \dots, \mathbf{x}_k), \mathbf{x}_{k+1}, \dots, \mathbf{x}_n).$$

By construction, note that $((T\mathbf{a}')_1, \dots, (T\mathbf{a}')_k) = E\mathbf{a}'_R = (0^{k-1}, 1)$. Next note that $((T\mathbf{a}')_{k+1}, \dots, (T\mathbf{a}')_n)$ are linearly independent over \mathbb{Q} . Letting $T \leftarrow T'T$ and $\mathbf{a}' \leftarrow T'\mathbf{a}'$, we have that $\mathbf{a}' = \lambda T\mathbf{a}$ satisfies the required form.

Given the above case analysis, we are left with the case where $\mathbf{a}'_R = (a'_1, \dots, a'_k) \in$

\mathbb{Q}^k , $\mathbf{a}'_I = (a'_{k+1}, \dots, a'_n) \in (\mathbb{R} \setminus \mathbb{Q})^{n-k}$ and where $(1, a'_{k+1}, \dots, a'_n)$ have a linear dependency over \mathbb{Q} . Now after an appropriate scaling of this dependency, we get numbers $c_0 \in \mathbb{Q}$, $\mathbf{c} \in \mathbb{Z}^{n-k} \setminus \{0\}$, $\gcd(c_1, \dots, c_{n-k}) = 1$, and where

$$\langle \mathbf{a}'_I, \mathbf{c} \rangle = c_0$$

Applying the Euclidean algorithm on \mathbf{c} , we get a unimodular matrix E such that

$$E\mathbf{c} = (\gcd(\mathbf{c}_1, \dots, \mathbf{c}_{n-k}), 0^{n-k-1}) = (1, 0^{n-k-1})$$

Let $\hat{\mathbf{a}} = E^{-t}\mathbf{a}'_I$. Note that E is unimodular iff E^{-t} is unimodular. We get that

$$\langle \mathbf{a}'_I, \mathbf{c} \rangle = c_0 \Rightarrow \langle E^{-t}\mathbf{a}'_I, E\mathbf{c} \rangle = c_0 \Rightarrow \hat{\mathbf{a}}_1 = c_0$$

Hence we see that $\hat{\mathbf{a}}_1 = c_0 \in \mathbb{Q}$. Let T' be the unimodular transformation defined by

$$T'(\mathbf{x}) = (\mathbf{x}_1, \dots, \mathbf{x}_k, E^{-t}(\mathbf{x}_{k+1}, \dots, \mathbf{x}_n))$$

Here T' is the identity on the first k coordinates, and acts like E^{-t} on the last $n - k$ coordinates. Note that $((T'\mathbf{a}')_1, \dots, (T'\mathbf{a}')_k) = (\mathbf{a}'_1, \dots, \mathbf{a}'_k) \in \mathbb{Q}^k$. Next $((T'\mathbf{a}')_{k+1}, \dots, (T'\mathbf{a}')_n) = E^{-t}\mathbf{a}'_I = \hat{\mathbf{a}}$, and $\hat{\mathbf{a}}_1 \in \mathbb{Q}$. Hence $T'\mathbf{a}'$ has at least one more rational coefficient than \mathbf{a}' . By repeating the above operation suitable number of times, we obtain a vector $\mathbf{a}' \in \mathbb{R}^n$ such that $\mathbf{a}'_1, \dots, \mathbf{a}'_k$ are rational and $(1, \mathbf{a}'_{k+1}, \dots, \mathbf{a}'_n)$ are linearly independent over \mathbb{Q} . By the previous analysis, there exists unimodular transformation T'' , $\lambda' \in \mathbb{Q}$ such that $\lambda'T''T'\mathbf{a}'$ satisfies the required form. Letting $T \leftarrow T''T'T$, $\lambda \leftarrow \lambda'\lambda$, and $\mathbf{a}' \leftarrow \lambda'T''T'\mathbf{a}'$, we get the desired result.

For proving the second part of the result, we first claim that $\mathcal{D}(a') = \mathcal{D}(a)$. To see this, note that

$$A\mathbf{a}' = 0 \Leftrightarrow A(\lambda T\mathbf{a}) = 0 \Leftrightarrow AT\mathbf{a} = 0$$

and

$$A\mathbf{a} = 0 \Leftrightarrow A\left(\frac{1}{\lambda}T^{-1}\mathbf{a}'\right) = 0 \Leftrightarrow AT^{-1}\mathbf{a}' = 0$$

since T is invertible and $\lambda \neq 0$. Since both AT, AT^{-1} are rational, this gives that $\mathcal{D}(\mathbf{a}') = \mathcal{D}(\mathbf{a})$ as needed. Hence we need only show that $\mathcal{D}(\mathbf{a}') = s + t$.

Take $\mathbf{y} \in \mathbb{Q}^n$ such that $\langle \mathbf{y}, \mathbf{a}' \rangle = 0$. Note that $\mathbf{a}' = (0^t, 1^s, \alpha_1, \dots, \alpha_r)$ where $(1, \alpha_1, \dots, \alpha_r)$ are linearly independent over \mathbb{Q} . If $s = 0$, then $\sum_{i=1}^r y_{t+i} \alpha_i = 0$. Since $\mathbf{y} \in \mathbb{Q}^n$, this gives a linear dependence of $(\alpha_1, \dots, \alpha_r)$ over \mathbb{Q} , and hence by assumption we must have that $y_{t+i} = 0$ for $1 \leq i \leq r$. Otherwise if $s = 1$, we get $y_{t+1} + \sum_{i=1}^r y_{t+i+1} \alpha_i = 0$, which gives a linear dependence of $(1, \alpha_1, \dots, \alpha_r)$ over \mathbb{Q} . Therefore $y_{t+i} = 0$ for $1 \leq i \leq t + 1$. Hence in both cases, we get that $y_{t+i} = 0$ for $1 \leq i \leq r + s$. Next note that for $\mathbf{y} \in \mathbb{Q}^t \times 0^{n-t}$, we have that $\langle \mathbf{y}, \mathbf{a}' \rangle = 0$ since $a'_1, \dots, a'_r = 0$ by assumption. By the previous observations, we obtain that

$$L := \{\mathbf{y} \in \mathbb{Q}^n : \langle \mathbf{y}, \mathbf{a}' \rangle = 0\} = \mathbb{Q}^t \times 0^{n-t} = \mathbb{Q}^t \times 0^{s+r}.$$

Now let $W \subseteq \mathbb{R}^n$ denote the linear subspace $W = \{\mathbf{x} \in \mathbb{R}^n : x_i = 0, 1 \leq i \leq t\}$. Note that $\mathbf{a}' \in W$, and hence $\mathcal{D}(\mathbf{a}') \leq \dim(W) = s + r$. Now take any $M = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{0}\}$, such that $\mathbf{a}' \in M$ and $A \in \mathbb{Q}^{m \times n}$. We claim that $W \subseteq M$. Let $a_1, \dots, a_m \in \mathbb{Q}^n$ denote the rows of A . Since $\mathbf{a}' \in M$, we have $\langle \mathbf{a}_i, \mathbf{a}' \rangle = 0 \forall i \in \{1, \dots, m\}$. Hence we must have that $\mathbf{a}_i \in L = \mathbb{Q}^t \times 0$. Since $W = 0^t \times \mathbb{R}^{s+r}$, we have that for all $\mathbf{x} \in W$, $\langle \mathbf{a}_i, \mathbf{x} \rangle = 0$, and hence $W \subseteq L$. Hence

$$\dim(L) \geq \dim(W) = s + r,$$

from which conclude that $\mathcal{D}(\mathbf{a}') = s + r$ as needed. \square

We now show that the points in $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$ can be separated using a finite number of CG cuts. We first give a rough sketch of the proof. We restrict to the case where $\text{aff}_I(H_{\mathbf{v}}) \neq \emptyset$. From here one can verify that any rational affine subspace contained in $\text{aff}(H_{\mathbf{v}})$ must also lie in $\text{aff}_I(H_{\mathbf{v}})$. Next we use Kronecker's theorem to build a finite set $C \subseteq \mathbb{Z}^n$, where each vector in C is at distance at most ϵ from some scaling of \mathbf{v} , and where \mathbf{v} can be expressed as a non-negative combination

of the vectors in C . By choosing ϵ and the scalings of \mathbf{v} appropriately, we can ensure that the CG cuts derived from C dominate the inequality $\langle \mathbf{v}, \mathbf{x} \rangle \leq h_K(\mathbf{v})$, i.e. $CG(K, C) \subseteq H_{\mathbf{v}}^{\leq}$. If $CG(K, C)$ lies in the interior of $H_{\mathbf{v}}^{\leq}(K)$, we have separated all of $H_{\mathbf{v}}$ (including $F_{\mathbf{v}} \setminus \text{aff}_I(H_{\mathbf{v}})$) and hence are done. Otherwise, $T := CG(K, C) \cap H_{\mathbf{v}}$ is a face of a rational polyhedron, and therefore $\text{aff}(T)$ is a rational affine subspace. Since $\text{aff}(T) \subseteq \text{aff}(H_{\mathbf{v}})$, as discussed above we obtain $T \subseteq \text{aff}(T) \subseteq \text{aff}_I(H_{\mathbf{v}})$ as required.

Proposition 3.3.9. *Let $K \subseteq \mathbb{R}^n$ be a compact convex set and $\mathbf{v} \in \mathbb{R}^n$. Then there exists $C \subseteq \mathbb{Z}^n$, $|C| \leq \mathcal{D}(\mathbf{v}) + 1$, such that*

$$CG(K, C) \subseteq H_{\mathbf{v}}^{\leq}(K) \quad \text{and} \quad CG(K, C) \cap H_{\mathbf{v}}(K) \subseteq \text{aff}_I(H_{\mathbf{v}}(K)).$$

Proof. By scaling \mathbf{v} by a positive scalar if necessary, we may assume that $h_K(\mathbf{v}) \in \{0, 1, -1\}$. Let T and λ denote the transformation and scaling promised for \mathbf{v} in Lemma 3.3.8. Note that $T^{-t}\{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{v}, \mathbf{x} \rangle = h_K(\mathbf{v})\} = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{v}, T^t \mathbf{x} \rangle = h_K(\mathbf{v})\} = \{\mathbf{x} \in \mathbb{R}^n : \langle \lambda T \mathbf{v}, \mathbf{x} \rangle = h_{T^{-t}K}(\lambda T \mathbf{v})\}$.

Now let $\mathbf{v}' = \lambda T \mathbf{v}$ and $b' = h_{T^{-t}K}(\lambda T \mathbf{v})$. By Lemma 3.3.7, it suffices to prove the statement for \mathbf{v}' and $K' = T^{-t}K$. Now \mathbf{v}' has the form (3.3.12) where $t, r \in \mathbb{Z}_+$, $s \in \{0, 1\}$, and $(1, \alpha_1, \dots, \alpha_r)$ are linearly independent over \mathbb{Q} . For convenience, let $k = s + t$, where we note that $\mathbf{v}'_{k+1}, \dots, \mathbf{v}'_{k+r} = (\alpha_1, \dots, \alpha_r)$.

Claim 1: Let $S = \{\mathbf{x} \in \mathbb{Z}^n : \langle \mathbf{v}', \mathbf{x} \rangle = b'\}$. Then S satisfies one of the following: (1) $S = \mathbb{Z}^t \times b' \times 0^r$: $s = 1, b' \in \mathbb{Z}$, (2) $S = \mathbb{Z}^t \times 0^r$: $s = 0, b' = 0$, (3) $S = \emptyset$: $s = 0, b' \neq 0$ or $s = 1, b' \notin \mathbb{Z}$.

Note that $b' = h_{T^{-t}K}(\lambda T \mathbf{v}) = \lambda h_K(\mathbf{v}) \in \{0, \pm\lambda\} \subseteq \mathbb{Q}$. We first see that

$$(s = 1) : b' = \langle \mathbf{v}', \mathbf{x} \rangle = \mathbf{x}_k + \sum_{i=1}^r \mathbf{x}_{k+i} \alpha_i, \quad (s = 0) : b' = \langle \mathbf{v}', \mathbf{x} \rangle = \sum_{i=1}^r \mathbf{x}_{k+i} \alpha_i.$$

Now if $\mathbf{x} \in S$, then

$$(s = 1) : (\mathbf{x}_k - b') + \sum_{i=1}^r \mathbf{x}_{k+i} \alpha_i = 0, \quad (s = 0) : (-b') + \sum_{i=1}^r \mathbf{x}_{k+i} \alpha_i = 0.$$

Since $b' \in \mathbb{Q}$, and $\mathbf{x} \in \mathbb{Z}^n$, in both cases the above equations give us a linear dependence of $(1, \alpha_1, \dots, \alpha_r)$ over \mathbb{Q} . Since by assumption $(1, \alpha_1, \dots, \alpha_r)$ are linearly independent over \mathbb{Q} , we have that

$$(s = 0, 1) : \mathbf{x}_{k+i} = 0, 1 \leq i \leq r \quad (s = 1) : \mathbf{x}_k = b' \quad (s = 0) : b' = 0.$$

If $s = 1$, then we must have that $b' \in \mathbb{Z}$, since $\mathbf{x}_k = b'$ and $\mathbf{x} \in \mathbb{Z}^n$. From this we immediately recover case (1). If $s = 0$, then the conditions $b' = 0$ and $\mathbf{x}_{k+i} = 0$, $1 \leq i \leq r$, verify case (2). If we are neither in case (1) or (2), then by the above analysis S must be empty, and so we are done.

Claim 2: Let $I = \{n\mathbf{v}' \pmod{1} : n \in N\}$. Then Theorem 3.3.6 implies that I is dense in $0^k \times [0, 1)^r$.

We first note that $\mathbf{v}'_1, \dots, \mathbf{v}'_k \in \mathbb{Z}$ and hence $\mathbf{v}'_1, \dots, \mathbf{v}'_k \equiv 0 \pmod{1}$. Next note that $(1, \alpha_1, \dots, \alpha_r)$ are linearly independent over \mathbb{Q} , and hence by Theorem 3.3.6 we have that $\{n(\alpha_1, \dots, \alpha_r) : n \in N\}$ is dense over $[0, 1)^r$. Putting the last two statements together immediately yields the claim.

Claim 3: There exists $\mathbf{a}_1, \dots, \mathbf{a}_{r+1} \subseteq \mathbb{Z}^n$ and $\lambda_1, \dots, \lambda_{r+1} \geq 0$ such that $\sum_{i=1}^{r+1} \lambda_i \mathbf{a}_i = \mathbf{v}'$ and $\sum_{i=1}^{r+1} \lambda_i [h'_K(\mathbf{a}_i)] \leq b'$, where the inequality can be made strict if $S = \emptyset$.

Since K' is compact, there exists $R > 0$ such that $K' \subseteq RB^n$. Take the subspace $W = 0^k \times \mathbb{R}^r$. Let $\mathbf{w}_1, \dots, \mathbf{w}_{r+1} \in W \cap S^{n-1}$, be any vectors such that for some $0 < \epsilon < 1$ we have $\sup_{1 \leq i \leq r+1} \langle \mathbf{w}_i, \mathbf{d} \rangle \geq \epsilon$ for all $\mathbf{d} \in S^{n-1} \cap W$ (e.g. $\mathbf{w}_1, \dots, \mathbf{w}_{r+1}$ are the vertices of a scaled isotropic r -dimensional simplex).

Case 1: $S \neq \emptyset$.

Let $a = \frac{1}{8} \min\{\frac{1}{R}, \epsilon\}$, and $b = \frac{1}{2}\epsilon a$. Now, for $1 \leq i \leq r + 1$ define $E_i = \{\mathbf{x} : \mathbf{x} \in a\mathbf{w}_i + b(B^n \cap W) \pmod{1}\}$. Since $W = 0^k \times \mathbb{R}^r$, note that $E_i \subseteq 0^k \times [0, 1)^r$.

By Claim 2 the set I is dense in $0^k \times [0, 1]^r$. Furthermore each set E_i has non-empty interior with respect to the subspace topology on $0^k \times [0, 1]^r$. Hence for all i , $1 \leq i \leq r+1$, we can find $n_i \in \mathbb{N}$ such that $n_i \mathbf{v}' \pmod{1} \in E_i$.

Now $n_i \mathbf{v}' \pmod{1} \in E_i$, implies that for some $\delta'_i \in E_i$, $n_i \mathbf{v}' - \delta'_i \in \mathbb{Z}^n$. Furthermore $\delta'_i \in E_i$ implies that there exists $\delta_i \in a\mathbf{w}_i + b(B^n \cap W)$ such that $\delta'_i - \delta_i \in \mathbb{Z}^n$. Hence $(n_i \mathbf{v}' - \delta'_i) + (\delta'_i - \delta_i) = n_i \mathbf{v}' - \delta_i \in \mathbb{Z}^n$. Let $a_i = n_i \mathbf{v}' - \delta_i$. Note that $\|\mathbf{a}_i - n_i \mathbf{v}'\| = \|\delta_i\| \leq a + b \leq 2a \leq 1/(4R)$. We claim that $\lfloor h_{K'}(\mathbf{a}_i) \rfloor \leq h_{K'}(n_i \mathbf{v}')$. First note that $h_{K'}(n_i \mathbf{v}') = n_i b'$. Since we assume that $S \neq \emptyset$, we must have that $b' \in \mathbb{Z}$ and hence $n_i b' \in \mathbb{Z}$. Now note that

$$\begin{aligned} h_{K'}(\mathbf{a}_i) &= h_{K'}((\mathbf{a}_i - n_i \mathbf{v}') + n_i \mathbf{v}') \leq h_{K'}(n_i \mathbf{v}') + h_{K'}(\mathbf{a}_i - n_i \mathbf{v}') \\ &= n_i b' + h_{K'}(-\delta_i) \\ &\leq n_i b' + h_{RB^n}(-\delta_i) \leq n_i b' + R\|\delta_i\| \leq n_i b' + R \left(\frac{1}{4R} \right) = n_i b' + \frac{1}{4}. \end{aligned}$$

Therefore we have that $\lfloor h_{K'}(\mathbf{a}_i) \rfloor \leq \lfloor n_i b' + \frac{1}{4} \rfloor = n_i b' = h_{K'}(n_i \mathbf{v}')$, since $n_i b' \in \mathbb{Z}$.

We claim that $\frac{a\epsilon}{4} B^n \cap W \subseteq \text{conv}\{\delta_1, \dots, \delta_{r+1}\}$. First note that by construction, $\text{conv}\{\delta_1, \dots, \delta_{r+1}\} \subseteq W$. Hence if the conclusion is false, then by the separator theorem there exists $\mathbf{d} \in W \cap S^{n-1}$ such that $h_{\frac{a\epsilon}{4} B^n \cap W}(\mathbf{d}) = \frac{a\epsilon}{4} > \sup_{1 \leq i \leq r+1} \langle \mathbf{d}, \delta_i \rangle$. For each i , $1 \leq i \leq r+1$, we write $\delta_i = a\mathbf{w}_i + b\mathbf{z}_i$ where $\|\mathbf{z}_i\| \leq 1$. Now note that

$$\begin{aligned} \sup_{1 \leq i \leq r+1} \langle \mathbf{d}, \delta_i \rangle &= \sup_{1 \leq i \leq r+1} \langle \mathbf{d}, a\mathbf{w}_i + b\mathbf{z}_i \rangle = \sup_{1 \leq i \leq r+1} a \langle \mathbf{d}, \mathbf{w}_i \rangle + b \langle \mathbf{d}, \mathbf{z}_i \rangle \\ &\geq \sup_{1 \leq i \leq r+1} a \langle \mathbf{d}, \mathbf{w}_i \rangle - b\|\mathbf{d}\|\|\mathbf{z}_i\| \geq a\epsilon - b = \frac{a\epsilon}{2} > \frac{a\epsilon}{4}, \end{aligned}$$

a contradiction. Hence there exists $\lambda_1, \dots, \lambda_{r+1} \geq 0$ and $\sum_{i=1}^{r+1} \lambda_i n_i = 1$ such that $\sum_{i=1}^{r+1} \lambda_i \delta_i = 0$.

Now we see that

$$\sum_{i=1}^{r+1} \lambda_i \mathbf{a}_i = \sum_{i=1}^{r+1} \lambda_i n_i \mathbf{v}' + \sum_{i=1}^{r+1} \lambda_i (\mathbf{a}_i - n_i \mathbf{v}') = \left(\sum_{i=1}^{r+1} \lambda_i n_i \right) \mathbf{v}' - \sum_{i=1}^{r+1} \lambda_i \delta_i = \left(\sum_{i=1}^{r+1} \lambda_i n_i \right) \mathbf{v}'.$$

Next note that

$$\sum_{i=1}^{r+1} \lambda_i \lfloor h_{K'}(\mathbf{a}_i) \rfloor \leq \sum_{i=1}^{r+1} \lambda_i h_{K'}(n_i \mathbf{v}') = h_{K'} \left(\left(\sum_{i=1}^{r+1} \lambda_i n_i \right) \mathbf{v}' \right).$$

Case 2: $S = \emptyset$. The proof here shall proceed very similarly to the one above, with the exception that we need to do some extra work to guarantee a strict inequality.

If $s = 0$, then since $S = \emptyset$ we must have that $b' \neq 0$. Let $\mathbf{v}^z = \frac{1}{|b'|} \mathbf{v}'$ and $b^z = \text{sign}(b')$, and $\mathbf{v}^f = \frac{1}{2|b'|} \mathbf{v}'$ and $b^f = \frac{1}{2} \text{sign}(b')$. Note that $h_{K'}(\mathbf{v}^z) = b^z \in \{\pm 1\}$ and $h_{K'}(\mathbf{v}^f) = b^f \in \{\pm 1/2\}$. Furthermore, since $b' \in \mathbb{Q}$, we see that

$$(1, \mathbf{v}_{k+1}^z, \dots, \mathbf{v}_{k+r}^z) = (1, \frac{1}{2|b'|} \alpha_1, \dots, \frac{1}{2|b'|} \alpha_r)$$

are still linearly independent over \mathbb{Q} , and that $\mathbf{v}_1^z, \dots, \mathbf{v}_k^z = \mathbf{v}_1', \dots, \mathbf{v}_k' = 0 \in \mathbb{Z}$.

Next if $s = 1$, then $b' \in \mathbb{Q} \setminus \mathbb{Z}$. Let $c_1 \in \mathbb{Z}$ denote the least positive integer such that $c_1 b' \in \mathbb{Z}$ and let $c_2 \in \mathbb{Z}$ denote the least positive integer such that $\frac{1}{3} \leq c_2 b' \pmod{1} \leq \frac{2}{3}$ (always exists since $b' \neq 0$). Let $\mathbf{v}^z = c_1 \mathbf{v}'$ and $b^z = c_1 b'$, and let $\mathbf{v}^f = c_2 \mathbf{v}'$ and $b^f = c_2 b'$. Again we have that $h_{K'}(\mathbf{v}^z) = b^z \in \mathbb{Z}$, and $h_{K'}(\mathbf{v}^f) = b^f$ (since $c_1, c_2 \geq 0$). Lastly, since $c_1, c_2 \in \mathbb{Z}$, we note that $\mathbf{v}_1^z, \dots, \mathbf{v}_{k-1}^z = \mathbf{v}_1^f, \dots, \mathbf{v}_{k-1}^f = 0 \in \mathbb{Z}$, $\mathbf{v}_k^z = c_1, \mathbf{v}_k^f = c_2 \in \mathbb{Z}$, and $(1, \mathbf{v}_{k+1}^z, \dots, \mathbf{v}_{k+r}^z) = (1, c_1 \alpha_1, \dots, c_1 \alpha_r)$ are still linearly independent over \mathbb{Q} .

Now let $I' = \{n \mathbf{v}^z \pmod{1} : n \in \mathbb{N}\}$. Using the proof of Claim 2, we see that I' is dense in $0^k \times [0, 1)^r$. Furthermore since $\mathbf{v}^f \pmod{1} \in 0^k \times [0, 1)^r$, we have that $I' + \mathbf{v}^f \pmod{1}$ is also dense in $0^k \times [0, 1)^r$. Note that $I' + \mathbf{v}^f \pmod{1} = \{(nc_1 + c_2) \mathbf{v}' \pmod{1} : n \in \mathbb{N}\}$.

Let $\mathbf{w}_1, \dots, \mathbf{w}_{l+1}, E_1, \dots, E_{l+1}$ be defined identically as in Case 1. Via the same density argument as in case 1, we may pick $n_i \in \mathbb{N}$, such that $(n_i c_1 + c_2) \mathbf{v}' \in E_i$. Again we define $\mathbf{a}_1, \dots, \mathbf{a}_{r+1}$ in exactly the same way as in Case 1. To conclude the proof of the claim, we need only show that $\lfloor h_{K'}(\mathbf{a}_i) \rfloor \leq \lfloor n_i b' + \frac{1}{4} \rfloor = n_i b' = h_{K'}(n_i \mathbf{v}')$ holds with a strict inequality in this case. The exact same argument gives us now that

$$h_{K'}(\mathbf{a}_i) \leq (n_i c_1 + c_2) b' + \frac{1}{4}. \quad (3.3.13)$$

Now $n_i c_1 b' = n_i b^z \in \mathbb{Z}$ and $\frac{1}{3} \leq c_2 b' \pmod{1} \leq \frac{2}{3}$. Therefore

$$\lfloor h_{K'}(\mathbf{a}_i) \rfloor < (n_i c_1 + c_2) b', \quad (3.3.14)$$

as needed.

Claim 4: Let $C = \{\mathbf{a}_i\}_{i=1}^{r+1}$ for the \mathbf{a}_i 's from Claim 3. Then $CG(K, C) \cap \{\mathbf{x} : \langle \mathbf{v}', \mathbf{x} \rangle = b'\} \subseteq \text{aff}(S)$.

If $S = \emptyset$, note that by the Claim 3, we have that

$$\sup\{\langle \mathbf{v}', \mathbf{x} \rangle : \mathbf{x} \in \mathbb{R}^n, \langle \mathbf{a}_i, \mathbf{x} \rangle \leq \lfloor h_{K'}(\mathbf{a}_i) \rfloor, 1 \leq i \leq r+1\} < b',$$

and hence $CG(K, C) \cap \{\mathbf{x} : \langle \mathbf{v}', \mathbf{x} \rangle = b'\} = \emptyset$ as needed.

If $S \neq \emptyset$, examine the set

$$P = \{\mathbf{x} : \langle \mathbf{v}', \mathbf{x} \rangle = b', \langle \mathbf{a}_i, \mathbf{x} \rangle \leq \lfloor h_{K'}(\mathbf{a}_i) \rfloor, 1 \leq i \leq l+1\}.$$

From the proof of Claim 3, we know that for each i , $1 \leq i \leq r+1$, we have $\lfloor h_{K'}(\mathbf{a}_i) \rfloor \leq h_{K'}(n_i \mathbf{v}') = n_i b'$ and hence $\langle n_i \mathbf{v}' - \mathbf{a}_i, \mathbf{x} \rangle = \langle \delta_i, \mathbf{x} \rangle \geq 0$, is a valid inequality for P . Now, from the proof of Claim 3, we have

$$\frac{a\epsilon}{4} B^n \cap W \subseteq \text{conv}\{\delta_1, \dots, \delta_{r+1}\}. \quad (3.3.15)$$

We claim that for all $H \subseteq \{1, \dots, r+1\}$, $|H| = r$, the set $\{\delta_i : i \in H\}$ is linearly independent. Assume not, then WLOG we may assume that $\delta_1, \dots, \delta_r$ are not linearly independent. Hence there exists $\mathbf{d} \in S^{n-1} \cap W$, such that $\langle \mathbf{d}, \delta_i \rangle = 0$ for all $1 \leq i \leq r$. Now by possibly switching \mathbf{d} to $-\mathbf{d}$, we may assume that $\langle \mathbf{d}, \delta_{r+1} \rangle \leq 0$. Hence we get that $\sup_{1 \leq i \leq r+1} \langle \mathbf{d}, \delta_i \rangle \leq 0$ in contradiction to (3.3.15).

Now let $\lambda_1, \dots, \lambda_{r+1} \geq 0$, $\sum_{i=1}^{r+1} \lambda_i n_i = 1$ be a combination such that $\sum_{i=1}^{r+1} \lambda_i \delta_i = 0$. Note that $\lambda_1, \dots, \lambda_{r+1}$ forms a linear dependency on $\delta_1, \dots, \delta_{r+1}$, and hence by the previous claim we must have that $\lambda_i > 0$ for all $1 \leq i \leq r+1$.

We claim for $P \subseteq W^\perp$. To see this, note that $0 = \langle \mathbf{x}, 0 \rangle = \langle \mathbf{x}, \sum_{i=1}^{r+1} \lambda_i \delta_i \rangle = \sum_{i=1}^{r+1} \lambda_i \langle \mathbf{x}, \delta_i \rangle$ for every $\mathbf{x} \in P$. Now since $\text{span}(\delta_1, \dots, \delta_{r+1}) = W$, we see that

$\langle \mathbf{x}, \delta_i \rangle = 0$ for all $1 \leq i \leq r + 1$ iff $\mathbf{x} \in W^\perp$. Hence if $\mathbf{x} \notin W^\perp$, then by the above equation and the fact that $\lambda_i > 0$ for all $i \in \{1, \dots, r + 1\}$, there exists $i, j \in \{1, \dots, r + 1\}$ such that $\langle \mathbf{x}, \delta_i \rangle > 0$ and $\langle \mathbf{x}, \delta_j \rangle < 0$. But then $\mathbf{x} \notin P$, since $\langle \mathbf{x}, \delta_j \rangle < 0$, a contradiction. Now $W = 0^k \times \mathbb{R}^r$, hence $W^\perp = \mathbb{R}^k \times 0^r$. To complete the proof we see that $P \subseteq \{\mathbf{x} : \mathbf{x} \in \mathbb{R}^k \times 0^r, \langle \mathbf{v}', \mathbf{x} \rangle = b'\} = \text{aff}(S)$. \square

3.3.3 Lifting the CG Closure of an Exposed Face of K

Proposition 3.3.10. *Let $K \subseteq \mathbb{R}^n$ be a compact convex set. Take $\mathbf{v} \in \mathbb{R}^n$. Assume that $CG(F_{\mathbf{v}}(K))$ is finitely generated. Then $\exists S \subseteq \mathbb{Z}^n$, $|S| < \infty$, such that $CG(K, S)$ is a polytope and*

$$CG(K, S) \cap H_{\mathbf{v}}(K) = CG(F_{\mathbf{v}}(K)) \quad (3.3.16)$$

$$CG(K, S) \subseteq H_{\mathbf{v}}^{\leq}. \quad (3.3.17)$$

Proof. The right to left containment in (3.3.16) is direct from $CG(F_{\mathbf{v}}(K)) \subseteq CG(K, S)$ as every CG cut for K is a CG cut for $F_{\mathbf{v}}(K)$. For the reverse containment and for (3.3.17) we proceed as follows.

Using Proposition 3.3.9 there exists $S_1 \subseteq \mathbb{Z}^n$ such that $CG(K, S_1) \cap H_{\mathbf{v}}(K) \subseteq \text{aff}_I(H_{\mathbf{v}}(K))$ and $CG(K, S_1) \subseteq \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{v}, \mathbf{x} \rangle \leq h_K(\mathbf{v})\}$. Next let $G \subseteq \mathbb{Z}^n$ be such that $CG(F_{\mathbf{v}}(K), G) = CG(F_{\mathbf{v}}(K))$. For each $\mathbf{w} \in G$, by Proposition 3.3.5 there exists $\mathbf{w}' \in \mathbb{Z}^n$ such that $CG(K, \mathbf{w}') \cap \text{aff}_I(H_{\mathbf{v}}(K)) \subseteq CG(F_{\mathbf{v}}(K), \mathbf{w}) \cap \text{aff}_I(H_{\mathbf{v}}(K))$. For each $\mathbf{w} \in G$, add \mathbf{w}' above to S_2 . Now note that

$$\begin{aligned} CG(K, S_1 \cup S_2) \cap H_{\mathbf{v}}(K) &= CG(K, S_1) \cap CG(K, S_2) \cap H_{\mathbf{v}}(K) \\ &\subseteq CG(K, S_2) \cap \text{aff}_I(H_{\mathbf{v}}(K)) \\ &= CG(F_{\mathbf{v}}(K), G) \cap \text{aff}(H_{\mathbf{v}}(K)) \subseteq CG(F_{\mathbf{v}}(K)). \end{aligned}$$

Now let $S_3 = \{\pm \mathbf{e}_i : 1 \leq i \leq n\}$. Note that since K is compact $CG(K, S_3)$ is a cuboid with bounded side lengths, and hence is a polytope. Letting $S = S_1 \cup S_2 \cup S_3$, yields the desired result. \square

We now obtain a generalization of the classical result known for rational polyhedra.

Corollary 3.3.11. *Let K be a compact convex set and let F be an exposed face of K , then we have that $CG(F) = CG(K) \cap F$.*

3.4 Approximation of the CG Closure

3.4.1 Approximation 1 of the CG Closure

In this section, we construct a first approximation of the CG closure of K . Under the assumption that the CG closure of every proper exposed face is finitely generated, we use a compactness argument to construct a finite set of CG cuts $S \subseteq \mathbb{Z}^n$ such that $CG(K, S) \subseteq K \cap \text{aff}_I(K)$. We use the following lemma to simplify the analysis of integral affine subspaces.

Lemma 3.4.1. *Take $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Then there exists $\lambda \in \mathbb{R}^m$ such that for $\mathbf{a}' = \lambda A$, $\mathbf{b}' = \lambda \mathbf{b}$, we have that $\{\mathbf{x} \in \mathbb{Z}^n : A\mathbf{x} = \mathbf{b}\} = \{\mathbf{x} \in \mathbb{Z}^n : \mathbf{a}'\mathbf{x} = \mathbf{b}'\}$.*

Proof. If $\{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\} = \emptyset$, then by Farka's Lemma there exists $\lambda \in \mathbb{R}^m$ such that $\lambda A = 0$ and $\lambda \mathbf{b} = 1$. Hence $\{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\} = \{\mathbf{x} \in \mathbb{R}^n : 0\mathbf{x} = 1\} = \emptyset$ as needed. We may therefore assume that $\{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\} \neq \emptyset$. Therefore we may also assume that the rows of the augmented matrix $[A | \mathbf{b}]$ are linearly independent.

Let $T = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_m)$, where $\mathbf{a}_1, \dots, \mathbf{a}_m$ are the rows of A . Define $r : T \rightarrow \mathbb{R}$ where for $\mathbf{w} \in T$ we let $r(\mathbf{w}) = \lambda \mathbf{b}$ for $\lambda \in \mathbb{R}^m$ where $\lambda A = \mathbf{w}$. Since the rows of A are linearly independent we obtain that r is well defined and is a linear operator. Let $S = \{\mathbf{x} \in \mathbb{Z}^n : A\mathbf{x} = \mathbf{b}\}$. For $\mathbf{z} \in \mathbb{Z}^n$, examine $T_{\mathbf{z}} = \{\mathbf{w} \in T : \langle \mathbf{w}, \mathbf{z} \rangle = r(\mathbf{w})\}$. By linearity of r , we see that $T_{\mathbf{z}}$ is a linear subspace of T . Note that for $\mathbf{z} \in \mathbb{Z}^n$, $T_{\mathbf{z}} = T$ iff $\mathbf{z} \in S$. Therefore $\forall \mathbf{z} \in \mathbb{Z}^n \setminus S$, we must have that $T_{\mathbf{z}} \neq T$, and hence $\dim(T_{\mathbf{z}}) \leq \dim(T) - 1$. Let m_T denote the Lebesgue measure on T . Since $\dim(T_{\mathbf{z}}) < \dim(T)$, we see that $m_T(T_{\mathbf{z}}) = 0$. Let $T' = \bigcup_{\mathbf{z} \in \mathbb{Z}^n \setminus S} T_{\mathbf{z}}$. Since $\mathbb{Z}^n \setminus S$ is countable, by the countable subadditivity of m_T we have that $m_T(T') \leq \sum_{\mathbf{z} \in \mathbb{Z}^n \setminus S} m_T(T_{\mathbf{z}}) = 0$. Since $m_T(T) = \infty$, we must have that $T \setminus T' \neq \emptyset$. Hence we may pick $\mathbf{a}' \in T \setminus T'$.

Letting $\mathbf{b}' = r(\mathbf{a}')$, we note that by construction there $\exists \lambda \in \mathbb{R}^m$ such that $\lambda A = \mathbf{a}'$ and $\lambda \mathbf{b} = \mathbf{b}'$. Hence for all $\mathbf{z} \in S$, $\lambda A \mathbf{z} = \lambda \mathbf{b} \Rightarrow \mathbf{a}' \mathbf{z} = \mathbf{b}'$. Now take $\mathbf{z} \in \mathbb{Z}^n \setminus S$. Now since $\mathbf{a}' \in T \setminus T'$, we have that $\mathbf{a}' \notin T_{\mathbf{z}}$. Hence $\mathbf{a}' \mathbf{z} \neq \mathbf{b}'$. Therefore we see that $\{\mathbf{x} \in \mathbb{Z}^n : \mathbf{a}' \mathbf{x} = \mathbf{b}'\} = \{\mathbf{x} \in \mathbb{Z}^n : A \mathbf{x} = \mathbf{b}\}$ as needed. \square

Proposition 3.4.2. *Let $\emptyset \neq K \subseteq \mathbb{R}^n$ be a compact convex set. If $CG(F_{\mathbf{v}}(K))$ is finitely generated for any proper exposed face $F_{\mathbf{v}}(K)$ then $\exists S \subseteq \mathbb{Z}^n$, $|S| < \infty$, such that $CG(K, S) \subseteq K \cap \text{aff}_I(K)$ and $CG(K, S)$ is a polytope.*

Proof. Let us express $\text{aff}(K)$ as $\{\mathbf{x} \in \mathbb{R}^n : A \mathbf{x} = \mathbf{b}\}$. Note that $\text{aff}(K) \neq \emptyset$ since $K \neq \emptyset$. By Lemma 3.4.1 there exists λ , $\mathbf{c} = \lambda A$ and $d = \lambda \mathbf{b}$, and such that $\text{aff}(K) \cap \mathbb{Z}^n = \{\mathbf{x} \in \mathbb{Z}^n : \langle \mathbf{c}, \mathbf{x} \rangle = \mathbf{b}\}$. Since $h_K(\mathbf{c}) = \mathbf{b}$ and $h_K(-\mathbf{c}) = -\mathbf{b}$, using Proposition 3.3.9 on \mathbf{c} and $-\mathbf{c}$, we can find $S_A \subseteq \mathbb{Z}^n$ such that $CG(K, S_A) \subseteq \text{aff}(\{\mathbf{x} \in \mathbb{Z}^n : \langle \mathbf{c}, \mathbf{x} \rangle = \mathbf{b}\}) = \text{aff}_I(K)$.

Express $\text{aff}(K)$ as $W + \mathbf{a}$, where $W \subseteq \mathbb{R}^n$ is a linear subspace and $\mathbf{a} \in \mathbb{R}^n$. Now take $\mathbf{v} \in W \cap S^{n-1}$. Note that $F_{\mathbf{v}}(K)$ is a proper exposed face and hence, by assumption, $CG(F_{\mathbf{v}}(K))$ is finitely generated. Hence by Proposition 3.3.10 there exists $S_{\mathbf{v}} \subseteq \mathbb{Z}^n$ such that $CG(K, S_{\mathbf{v}})$ is a polytope, $CG(K, S_{\mathbf{v}}) \cap H_{\mathbf{v}}(K) = CG(F_{\mathbf{v}}(K))$ and $CG(K, S_{\mathbf{v}}) \subseteq \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{v} \rangle \leq h_K(\mathbf{v})\}$. Let $K_{\mathbf{v}} = CG(K, S_{\mathbf{v}})$, then we have the following claim.

Claim: \exists open neighborhood $N_{\mathbf{v}}$ of \mathbf{v} in $W \cap S^{n-1}$ such that $\mathbf{v}' \in N_{\mathbf{v}} \Rightarrow h_{K_{\mathbf{v}}}(\mathbf{v}') \leq h_K(\mathbf{v}')$.

Since $K_{\mathbf{v}}$ is a polytope, there exists $C \subseteq \mathbb{R}^n$, $|C| < \infty$, such that $K_{\mathbf{v}} = \text{conv}(C)$. Then note that $h_{K_{\mathbf{v}}}(\mathbf{w}) = \sup_{\mathbf{c} \in C} \langle \mathbf{c}, \mathbf{w} \rangle$. Now let $H = \{\mathbf{c} : h_K(\mathbf{v}) = \langle \mathbf{v}, \mathbf{c} \rangle, \mathbf{c} \in C\}$. By construction, we have that $\text{conv}(H) = CG(F_{\mathbf{v}}(K))$.

First assume that $CG(F_{\mathbf{v}}(K)) = \emptyset$. Then $H = \emptyset$, and hence $h_{K_{\mathbf{v}}}(\mathbf{v}) < h_K(\mathbf{v})$. Since $K_{\mathbf{v}}, K$ are compact convex sets, we have that $h_{K_{\mathbf{v}}}, h_K$ are both continuous functions on \mathbb{R}^n and hence $h_K - h_{K_{\mathbf{v}}}$ is continuous. Therefore there exists $\epsilon > 0$ such

that $h_{K_{\mathbf{v}}}(\mathbf{v}') < h_K(\mathbf{v}')$ for $\|\mathbf{v} - \mathbf{v}'\| \leq \epsilon$ as needed.

Now assume that $CG(F_{\mathbf{v}}(K)) \neq \emptyset$. Let $R = \max_{\mathbf{c} \in C} \|\mathbf{c}\|$, and let

$$\delta = h_K(\mathbf{v}) - \sup\{\langle \mathbf{v}, \mathbf{c} \rangle : \mathbf{c} \in C \setminus H\}.$$

Now let $\epsilon = \frac{\delta}{2R}$. Now take any \mathbf{v}' such that $\|\mathbf{v}' - \mathbf{v}\| < \epsilon$. Now for all $\mathbf{c} \in H$, we have that

$$\begin{aligned} \langle \mathbf{c}, \mathbf{v}' \rangle &= \langle \mathbf{c}, \mathbf{v} \rangle + \langle \mathbf{c}, \mathbf{v}' - \mathbf{v} \rangle = h_K(\mathbf{v}) + \langle \mathbf{c}, \mathbf{v}' - \mathbf{v} \rangle \geq h_K(\mathbf{v}) - \|\mathbf{c}\| \|\mathbf{v}' - \mathbf{v}\| \\ &> h_K(\mathbf{v}) - R \frac{\delta}{2R} = h_K(\mathbf{v}) - \frac{\delta}{2}, \end{aligned}$$

and that for all $\mathbf{c} \in C \setminus H$, we have that

$$\begin{aligned} \langle \mathbf{c}, \mathbf{v}' \rangle &= \langle \mathbf{c}, \mathbf{v} \rangle + \langle \mathbf{c}, \mathbf{v}' - \mathbf{v} \rangle \leq h_K(\mathbf{v}) - \delta + \langle \mathbf{c}, \mathbf{v}' - \mathbf{v} \rangle \leq h_K(\mathbf{v}) - \delta + \|\mathbf{c}\| \|\mathbf{v}' - \mathbf{v}\| \\ &< h_K(\mathbf{v}) - \delta + \frac{\delta}{2} = h_K(\mathbf{v}) - \frac{\delta}{2}. \end{aligned}$$

Therefore we have that $\langle c, v' \rangle > \langle c', v' \rangle$ for all $c \in H$, $c' \in C \setminus H$ and hence

$$h_{K_{\mathbf{v}}}(\mathbf{v}') = \sup_{c \in C} \langle c, \mathbf{v}' \rangle = \sup_{c \in H} \langle c, \mathbf{v}' \rangle = h_{CG(F_{\mathbf{v}}(K))}(\mathbf{v}') \leq h_K(\mathbf{v}'), \quad (3.4.1)$$

since $CG(F_{\mathbf{v}}(K)) \subseteq F_{\mathbf{v}}(K) \subseteq K$. The statement thus holds by letting $N_{\mathbf{v}} = \{\mathbf{v}' \in S^{n-1} : \|\mathbf{v}' - \mathbf{v}\| \leq \epsilon\}$.

Note that $\{N_{\mathbf{v}} : \mathbf{v} \in W \cap S^{n-1}\}$ forms an open cover of $W \cap S^{n-1}$, and since $W \cap S^{n-1}$ is compact, there exists a finite subcover $N_{\mathbf{v}_1}, \dots, N_{\mathbf{v}_k}$ such that $\bigcup_{i=1}^k N_{\mathbf{v}_i} = W \cap S^{n-1}$. Now let $S = S_A \cup \bigcup_{i=1}^k S_{\mathbf{v}_i}$. We claim that $CG(K, S) \subseteq K$. Assume not, then there exists $\mathbf{x} \in CG(K, S) \setminus K$. Since $CG(K, S) \subseteq CG(K, S_A) \subseteq W + \mathbf{a}$ and $K \subseteq W + \mathbf{a}$, by the separator theorem there exists $\mathbf{w} \in W \cap S^{n-1}$ such that $h_K(\mathbf{w}) = \sup_{\mathbf{y} \in K} \langle \mathbf{y}, \mathbf{w} \rangle < \langle \mathbf{x}, \mathbf{w} \rangle \leq h_{CG(K, S)}(\mathbf{w})$. Since $\mathbf{w} \in W \cap S^{n-1}$, there exists i , $1 \leq i \leq k$, such that $\mathbf{w} \in N_{\mathbf{v}_i}$. Note then we obtain that $h_{CG(K, S)}(\mathbf{w}) \leq h_{CG(K, S_{\mathbf{v}_i})}(\mathbf{w}) = h_{K_{\mathbf{v}_i}}(\mathbf{w}) \leq h_K(\mathbf{w})$, a contradiction. Hence $CG(K, S) \subseteq K$ as claimed. $CG(K, S)$ is a polytope because it is the intersection of polyhedra of which at least one is a polytope. \square

3.4.2 Approximation 2 of the CG Closure

In this section, we augment the first approximation of the $CG(K)$ with a finite number of extra CG cuts so that this second approximation matches $CG(K)$ on the relative boundary of K .

To achieve this, we observe that our first approximation of $CG(K)$ is polyhedral and contained in K , and hence its intersection with the relative boundary of K is contained in the union of a finite number of proper exposed faces of K . Therefore, by applying Proposition 3.3.10 to each such face (i.e. adding their lifted CG closure), we can match $CG(K)$ on the relative boundary as required. The following lemma makes precise the previous statements.

Lemma 3.4.3. *Let $K \subseteq \mathbb{R}^n$ be a convex set and $P \subseteq K$ be a polytope. Then there exists $F_{\mathbf{v}_1}, \dots, F_{\mathbf{v}_k} \subseteq K$, proper exposed faces of K , such that $P \cap \text{relbd}(K) \subseteq \bigcup_{i=1}^k F_{\mathbf{v}_i}$*

Proof. Let $\mathcal{F} = \{F : F \subseteq P, F \text{ a face of } P, \text{relint}(F) \cap \text{relbd}(K) \neq \emptyset\}$. Since P is polytope, note that the total number of faces of P is finite, and hence $|\mathcal{F}| < \infty$. We claim that

$$P \cap \text{relbd}(K) \subseteq \bigcup_{F \in \mathcal{F}} F. \quad (3.4.2)$$

Take $\mathbf{x} \in P \cap \text{relbd}(K)$. Let $F_{\mathbf{x}}$ denote the minimal face of P containing \mathbf{x} (note that P is a face of itself). By minimality of $F_{\mathbf{x}}$, we have that $\mathbf{x} \in \text{relint}(F_{\mathbf{x}})$. Since $\mathbf{x} \in \text{relbd}(K)$, we have that $F_{\mathbf{x}} \in \mathcal{F}$, as needed.

Take $F \in \mathcal{F}$. We claim that there exists $H_F \subseteq K$, H_F a proper exposed face of K , such that $F \subseteq H_F$. Take $\mathbf{x} \in \text{relint}(F) \cap \text{relbd}(K)$. Let $\text{aff}(K) = W + \mathbf{a}$, where W is a linear subspace and $\mathbf{a} \in \mathbb{R}^n$. Since $\mathbf{x} \notin \text{relint}(K)$, by the separator theorem, there exists $\mathbf{v} \in W \cap S^{n-1}$ such that $h_K(\mathbf{v}) = \langle \mathbf{x}, \mathbf{v} \rangle$. Let $H_F = F_{\mathbf{v}}(K)$. Note that since $\mathbf{v} \in W \cap S^{n-1}$, $F_{\mathbf{v}}(K)$ is a proper exposed face of K . We claim that $F \subseteq H_F$. Since F is a polytope, we have that $F = \text{conv}(\text{ext}(F))$. Write $\text{ext}(F) = \{\mathbf{c}_1, \dots, \mathbf{c}_k\}$. Now

since $\mathbf{x} \in \text{relint}(F)$, there exists $\lambda_1, \dots, \lambda_k > 0$, $\sum_{i=1}^k \lambda_i = 1$, such that $\sum_{i=1}^k \lambda_i \mathbf{c}_i = \mathbf{x}$. Now since $\mathbf{c}_i \in K$, we have that $\langle \mathbf{c}_i, \mathbf{v} \rangle \leq h_K(\mathbf{v})$. Therefore, we note that

$$\langle \mathbf{x}, \mathbf{v} \rangle = \left\langle \sum_{i=1}^k \lambda_i \mathbf{c}_i, \mathbf{v} \right\rangle = \sum_{i=1}^k \lambda_i \langle \mathbf{c}_i, \mathbf{v} \rangle \leq \sum_{i=1}^k \lambda_i h_K(\mathbf{v}) = h_K(\mathbf{v}) \quad (3.4.3)$$

Since $\langle \mathbf{x}, \mathbf{v} \rangle = h_K(\mathbf{v})$, we must have equality throughout. To maintain equality, since $\lambda_i > 0$, $1 \leq i \leq k$, we must have that $\langle \mathbf{c}_i, \mathbf{v} \rangle = h_K(\mathbf{v})$, $1 \leq i \leq k$. Therefore $\mathbf{c}_i \in H_F$, $1 \leq i \leq k$, and hence $F = \text{conv}(\mathbf{c}_1, \dots, \mathbf{c}_k) \subseteq H_F$, as needed.

To conclude the proof, we note that the set $\{H_F : F \in \mathcal{F}\}$ satisfies the conditions of the lemma. □

Proposition 3.4.4. *Let $K \subseteq \mathbb{R}^n$ be a compact convex set. If $CG(F_{\mathbf{v}})$ is finitely generated for any proper exposed face $F_{\mathbf{v}}$ then $\exists S \subseteq \mathbb{Z}^n$, $|S| < \infty$, such that*

$$CG(K, S) \subseteq K \cap \text{aff}_I(K) \quad (3.4.4)$$

$$CG(K, S) \cap \text{relbd}(K) = CG(K) \cap \text{relbd}(K) \quad (3.4.5)$$

Proof. By Proposition 3.4.2, there exists $S_I \subseteq \mathbb{Z}^n$, $|S_I| < \infty$, such that $CG(K, S_I) \subseteq K \cap \text{aff}_I(K)$ and $CG(K, S_I)$ is a polytope. Since $CG(K, S_I) \subseteq K$ is a polytope, let $F_{\mathbf{v}_1}, \dots, F_{\mathbf{v}_k}$ be the proper exposed faces of K given by Lemma 3.4.3. By Proposition 3.3.10, there exists $S_i \subseteq \mathbb{Z}^n$, $|S_i| < \infty$, such that $CG(K, S_i) \cap H_{\mathbf{v}_i} = CG(F_{\mathbf{v}_i})$. Let $S = S_I \cup \bigcup_{i=1}^k S_i$. We claim that $CG(K, S) \cap \text{relbd}(K) \subseteq CG(K) \cap \text{relbd}(K)$. For this note that $\mathbf{x} \in CG(K, S) \cap \text{relbd}(K)$ implies $\mathbf{x} \in CG(K, S_I) \cap \text{relbd}(K)$, and hence there exists i , $1 \leq i \leq k$, such that $\mathbf{x} \in F_{\mathbf{v}_i}$. Then $\mathbf{x} \in CG(K, S) \cap H_{\mathbf{v}_i} \subseteq CG(K, S_i) \cap H_{\mathbf{v}_i} = CG(F_{\mathbf{v}_i}) \subseteq CG(K) \cap \text{relbd}(K)$. The reverse inclusion is direct. □

3.5 Proof of Theorem

Finally, we have all the ingredients to prove the main result of the Chapter. The proof is by induction on the dimension of K . Trivially, the result holds for zero

dimensional convex bodies. Now using the induction hypothesis, we can construct the second approximation of $CG(K)$ using Proposition 3.4.4 (since it assumes that the CG closure of every exposed face is finitely generated). Lastly, we observe that any CG cut for K not dominated by those already considered in the second approximation of $CG(K)$ must separate a vertex of this approximation lying in the relative interior of K . From here, it is not difficult to show that only a finite number of such cuts exists, thereby proving the polyhedrality of $CG(K)$. The proof here is similar to the one used for strictly convex sets, with the additional technicality that here $\text{aff}(K)$ may be irrational.

Theorem 3.5.1. *Let $K \subseteq \mathbb{R}^n$ be a non-empty compact convex set. Then $CG(K)$ is finitely generated.*

Proof. We proceed by induction on the affine dimension of K . For the base case, $\dim(\text{aff}(K)) = 0$, i.e. $K = \{\mathbf{x}\}$ is a single point. Here it is easy to see that setting $S = \{\pm \mathbf{e}_i : i \in \{1, \dots, n\}\}$, we get that $CG(K, S) = CG(K)$. The base case thus holds.

Now for the inductive step let $0 \leq k < n$ let K be a compact convex set where $\dim(\text{aff}(K)) = k + 1$ and assume the result holds for sets of lower dimension. By the induction hypothesis, we know that $CG(F_{\mathbf{v}})$ is finitely generated for every proper exposed face $F_{\mathbf{v}}$ of K , since $\dim(F_{\mathbf{v}}) \leq k$. By Proposition 3.4.4, there exists a set $S \subseteq \mathbb{Z}^n$, $|S| < \infty$, such that (3.4.4) and (3.4.5) hold. If $CG(K, S) = \emptyset$, then we are done. So assume that $CG(K, S) \neq \emptyset$. Let $A = \text{aff}_I(K)$. Since $CG(K, S) \neq \emptyset$, we have that $A \neq \emptyset$ (by (3.4.4)), and so we may pick $\mathbf{t} \in A \cap \mathbb{Z}^n$. Note that $A - \mathbf{t} = W$, where W is a linear subspace of \mathbb{R}^n satisfying $W = \text{span}(W \cap \mathbb{Z}^n)$. Let $L = W \cap \mathbb{Z}^n$. Since $\mathbf{t} \in \mathbb{Z}^n$, we easily see that $CG(K - \mathbf{t}, T) = CG(K, T) - \mathbf{t}$ for all $T \subseteq \mathbb{Z}^n$. Therefore $CG(K)$ is finitely generated iff $CG(K - \mathbf{t})$ is. Hence replacing K by $K - \mathbf{t}$, we may assume that $\text{aff}_I(K) = W$.

Let π_W denote the orthogonal projection onto W . Note that for all $\mathbf{x} \in W$,

and $\mathbf{z} \in \mathbb{Z}^n$, we have that $\langle \mathbf{z}, \mathbf{x} \rangle = \langle \pi_W(\mathbf{z}), \mathbf{x} \rangle$. Now since $CG(K, S) \subseteq K \cap W$, we see that for all $\mathbf{z} \in \mathbb{Z}^n$, $CG(K, S \cup \{\mathbf{z}\}) = CG(K, S) \cap \{\mathbf{x} : \langle \mathbf{z}, \mathbf{x} \rangle \leq \lfloor h_K(\mathbf{z}) \rfloor\} = CG(K, S) \cap \{\mathbf{x} : \langle \pi_W(\mathbf{z}), \mathbf{x} \rangle \leq \lfloor h_K(\mathbf{z}) \rfloor\}$. Let $L^* = \pi_W(\mathbb{Z}^n)$. Since W is a rational subspace, we have that L^* is full dimensional lattice in W . Now fix an element of $\mathbf{w} \in L^*$ and examine $V_{\mathbf{w}} := \{\lfloor h_K(\mathbf{z}) \rfloor : \pi_W(\mathbf{z}) = \mathbf{w}, \mathbf{z} \in \mathbb{Z}^n\}$. Note that $V_{\mathbf{w}} \subseteq \mathbb{Z}$. We claim that $\inf(V_{\mathbf{w}}) \geq -\infty$. To see this, note that

$$\begin{aligned} \inf\{\lfloor h_K(\mathbf{z}) \rfloor : \pi_W(\mathbf{z}) = \mathbf{w}, \mathbf{z} \in \mathbb{Z}^n\} &\geq \inf\{\lfloor h_{K \cap W}(\mathbf{z}) \rfloor : \pi_W(\mathbf{z}) = \mathbf{w}, \mathbf{z} \in \mathbb{Z}^n\} \\ &= \inf\{\lfloor h_{K \cap W}(\pi_W(\mathbf{z})) \rfloor : \pi_W(\mathbf{z}) = \mathbf{w}, \mathbf{z} \in \mathbb{Z}^n\} \\ &= \lfloor h_{K \cap W}(\mathbf{w}) \rfloor > -\infty. \end{aligned}$$

Now since $V_{\mathbf{w}}$ is a lower bounded set of integers, there exists $\mathbf{z}_{\mathbf{w}} \in \pi_W^{-1}(\mathbf{w}) \cap \mathbb{Z}^n$ such that $\inf(V_{\mathbf{w}}) = \lfloor h_K(\mathbf{z}_{\mathbf{w}}) \rfloor$. From the above reasoning, we see that $CG(K, S \cup \pi_W^{-1}(\mathbf{z}) \cap \mathbb{Z}^n) = CG(K, S \cup \{\mathbf{z}_{\mathbf{w}}\})$. Now examine the set

$$C = \{\mathbf{w} : \mathbf{w} \in L^*, CG(K, S \cup \{\mathbf{z}_{\mathbf{w}}\}) \subsetneq CG(K, S)\}.$$

Here we get that

$$CG(K) = CG(K, S \cup \mathbb{Z}^n) = CG(K, S \cup \{\mathbf{z}_{\mathbf{w}} : \mathbf{w} \in L^*\}) = CG(K, S \cup \{\mathbf{z}_{\mathbf{w}} : \mathbf{w} \in C\}).$$

From the above equation, if we show that $|C| < \infty$, then $CG(K)$ is finitely generated. To do this, we will show that there exists $R > 0$, such that $C \subseteq RB^n$, and hence $C \subseteq L^* \cap RB^n$. Since L^* is a lattice, $|L^* \cap RB^n| < \infty$ for any fixed R , and so we are done.

Now let $P = CG(K, S)$. Since P is a polytope, we have that $P = \text{conv}(\text{ext}(P))$. Let $I = \text{ext}(P) \cap \text{relint}(K)$, and let $B = \text{ext}(P) \cap \text{relbd}(K)$. Hence $\text{ext}(P) = I \cup B$. By assumption on $CG(K, S)$, we know that for all $\mathbf{v} \in B$, we have that $\mathbf{v} \in CG(K)$. Hence for all $\mathbf{z} \in \mathbb{Z}^n$, we must have that $\langle \mathbf{z}, \mathbf{v} \rangle \leq \lfloor h_K(\mathbf{z}) \rfloor$ for all $\mathbf{v} \in B$. Now assume that for some $\mathbf{z} \in \mathbb{Z}^n$, $CG(K, S \cup \{\mathbf{z}\}) \subsetneq CG(K, S) = P$. We claim that

$\langle \mathbf{z}, \mathbf{v} \rangle > \lfloor h_K(\mathbf{z}) \rfloor$ for some $\mathbf{v} \in I$. If not, then $\langle \mathbf{v}, \mathbf{z} \rangle \leq \lfloor h_K(\mathbf{z}) \rfloor$ for all $\mathbf{v} \in \text{ext}(P)$, and hence $CG(K, S \cup \{\mathbf{z}\}) = CG(K, S)$, a contradiction. Hence such a $\mathbf{v} \in I$ must exist.

For $\mathbf{z} \in \mathbb{Z}^n$, note that $h_K(\mathbf{z}) \geq h_{K \cap W}(\mathbf{z}) = h_{K \cap W}(\pi_W(\mathbf{z}))$. Hence $\langle \mathbf{z}, \mathbf{v} \rangle > \lfloor h_K(\mathbf{z}) \rfloor$ for $\mathbf{v} \in I$ only if $\langle \pi_W(\mathbf{z}), \mathbf{v} \rangle = \langle \mathbf{z}, \mathbf{v} \rangle > \lfloor h_{K \cap W}(\pi_W(\mathbf{z})) \rfloor$. Let $C' := \{\mathbf{w} \in L^* : \exists \mathbf{v} \in I, \langle \mathbf{v}, \mathbf{w} \rangle > \lfloor h_{K \cap W}(\mathbf{w}) \rfloor\}$. From the previous discussion, we see that $C \subseteq C'$.

Since $I \subseteq \text{relint}(K) \cap W = \text{relint}(K \cap W)$ we have

$$\delta_{\mathbf{v}} = \sup\{r \geq 0 : rB^n \cap W + \mathbf{v} \subseteq K \cap W\} > 0$$

for all $\mathbf{v} \in I$. Let $\delta = \inf_{\mathbf{v} \in I} \delta_{\mathbf{v}}$. Since $|I| < \infty$, we see that $\delta > 0$. Now let $R = \frac{1}{\delta}$. Take $\mathbf{w} \in L^*$, $\|\mathbf{w}\| \geq R$. Note that $\forall \mathbf{v} \in I$,

$$\lfloor h_{K \cap W}(\mathbf{w}) \rfloor \geq h_{K \cap W}(\mathbf{w}) - 1 \geq h_{(\mathbf{v} + \delta B^n) \cap W}(\mathbf{w}) - 1 = \langle \mathbf{v}, \mathbf{w} \rangle + \delta \|\mathbf{w}\| - 1 \geq \langle \mathbf{v}, \mathbf{w} \rangle.$$

Hence $\mathbf{w} \notin C'$. Therefore $C \subseteq C' \subseteq RB^n$ and $CG(K)$ is finitely generated. \square

3.6 Conclusion

The need to solve non-linear IP models has rapidly expanded over the last years, and this trend is likely to continue for the foreseeable future. Given the usefulness of cutting planes in ILP, an important research direction is to understand the properties of cutting plane closures in the non-linear setting. In this Chapter, we have made significant progress in extending the study of the CG closure to this setting. Our main result was to show that the CG closure of any compact set is polyhedral. In the process of proving this, we believe we have developed useful tools for a more general study of cutting plane closures in this setting. As a consequence of our result, we also resolve an open question of Schrijver [119], who asked whether the CG closure of irrational polytopes is polyhedral.

Future Research. A first avenue for future research is to understand whether our CG closure result can be extended to other classic cutting plane closures. Perhaps, a first natural candidate, is the split closure. A split disjunction is indexed by an integer vector $\mathbf{y} \in \mathbb{Z}^n$, and an integer $\pi_0 \in \mathbb{Z}$. For a convex set $K \subseteq \mathbb{R}^n$, the split disjunction induced by \mathbf{y} and π_0 is

$$K^{\mathbf{y}, \pi_0} = \text{conv}\{K \cap H_{\mathbf{y}, \pi_0}^{\leq}, K \cap H_{\mathbf{y}, \pi_0+1}^{\geq}\}$$

The split closure of K is $SC(K) = \bigcap_{\mathbf{y} \in \mathbb{Z}^n, \pi_0 \in \mathbb{Z}} K^{\mathbf{y}, \pi_0}$. In [35], we show that the split closure of an ellipsoid not be polyhedral, so we cannot hope for a general polyhedrality result as is satisfied by the CG closure. However, we show that for a strictly convex body (a body whose boundary does not contain lines), the split closure is in fact finitely generated, i.e. there exists a finite number of split disjunctions that generate $SC(K)$. An interesting open question is whether the same is true for general compact sets. Another direction of interest, is to extend our polyhedrality result for the CG closure to unbounded convex sets. We know that the result cannot extend to all unbounded convex sets, as the CG closure of an irrational cone is not polyhedral. However, the CG closure of any rational polyhedron is indeed polyhedral. Therefore, exactly classifying the convex sets for which the CG closure is polyhedral remains an interesting open question.

CHAPTER IV

THE M-ELLIPSOID AND VOLUME ESTIMATION

The M-Ellipsoid is a fundamental construct in asymptotic convex geometry. An M-Ellipsoid for an n -dimensional convex body K is an ellipsoid E whose covering estimates with respect to K are single exponential, i.e. $2^{O(n)}$ translates of E suffices to cover K and vice versa. The existence of the M-Ellipsoid is a fundamental result due to Milman [92], which has lead to many fundamental discoveries in convex geometry.

In this Chapter, we give algorithms for constructing an M-Ellipsoid for any convex body, and provide an application to volume estimation in the oracle model. In particular, we give a nearly optimal deterministic algorithm for estimating the volume of any symmetric convex body. We provide further applications of the M-Ellipsoid to classical lattice problems and the integer programming problem in Chapters 5 and 7.

This Chapter is based on work from the papers [36] (joint with Chris Peikert and Santosh Vempala) and [32] (joint with Santosh Vempala).

4.1 Introduction

Ellipsoids have traditionally played an important role in the study of convex bodies. The classical Löwner-John ellipsoids, for instance, is the starting point for many interesting studies. To recall John's theorem, for any convex body K in \mathbb{R}^n , there is an ellipsoid E and center $\mathbf{x} \in \mathbb{R}^n$ such that

$$\mathbf{x} + E \subseteq K \subseteq \mathbf{x} + nE.$$

In fact, this bound is achieved by the *maximum volume* ellipsoid contained in K .

Ellipsoids have also been critical to the design and analysis of efficient algorithms. The most notable example is the ellipsoid algorithm [122, 127] for linear [74] and

convex optimization [56], which represents a frontier of polynomial-time solvability. For the basic problems of sampling and integration in high dimensions, the *inertial* ellipsoid defined by the covariance matrix of a distribution has played an important role in the development of efficient algorithms [69, 88, 126]. This ellipsoid also achieves the bounds of John’s theorem for general convex bodies (for centrally-symmetric convex bodies, the max-volume ellipsoid achieves the best possible sandwiching ratio of \sqrt{n} while the inertial ellipsoid could still have a ratio of n).

Another ellipsoid that has played a critical role in the development of modern convex geometry is the M-Ellipsoid (Milman’s ellipsoid). This object was introduced by Milman as a tool to prove fundamental inequalities in convex geometry such as the Bourgain-Milman and reverse Brunn-Minkowski inequality (see e.g., Chapter 7 of [106]). An M-Ellipsoid E of a convex body K has small *covering numbers* with respect to K . As shown by Milman, every convex body K has an ellipsoid E for which the number of translates of E needed to cover K and vice versa is bounded by $2^{O(n)}$ (i.e. $N(K, E)N(E, K) = 2^{O(n)}$). This is the best possible bound up to a constant in the exponent. In contrast, the John ellipsoid can have this covering bound as high as $n^{\Omega(n)}$. There are now multiple proofs of existence of the M-Ellipsoid: the original construction due to Milman [92], multiple ones by Pisier [106], and most recently, by Klartag [77].

The complexity of computing these ellipsoids is important for the applications mentioned above, but is also interesting to study for its own sake. John ellipsoids are NP-hard to compute, but their worst case sandwiching bounds can be approximated deterministically to within $O(\sqrt{n})$ in polynomial time via the ellipsoid method [56]. Inertial ellipsoids can be approximated to arbitrary accuracy by random sampling in polynomial time [69]. The associated question for M-Ellipsoids however has, to the best of our knowledge, not been considered previously.

Here we consider the task of constructing the M-Ellipsoid and explore its application to volume estimation. The extent to which randomness is essential for efficiency in computation is a very interesting and important question in computational complexity. Here we use the M-Ellipsoid to improve the deterministic complexity of volume estimation, a problem where a strong separation between randomized and deterministic complexity is known in the oracle model [51, 44]. In Chapter 5, we provide further applications of the M-Ellipsoid to classical lattice problems, namely the Shortest Vector Problem (SVP) and Closest Vector Problem (CVP), and in Chapter 7 we give its application to the Integer Programming Problem (IP).

4.1.1 Results

The main results of this chapter are two algorithms for computing M-Ellipsoids for arbitrary convex bodies in the oracle model. The first algorithm is based on an M-Ellipsoid construction of Klartag [77] and runs in randomized polynomial time. The second, based on Milman’s original construction [92], runs in deterministic $2^{O(n)}$ time and uses polynomial space. Lastly, we show that Milman’s construction can be modified to give a nearly optimal algorithm for deterministically estimating the volume of a symmetric convex body. A crucial tool we develop for this purpose, which will have many applications in the next chapters, is a near optimal algorithm for computing a cover of a convex body by an ellipsoid.

Here the input convex bodies will be specified by well guaranteed membership oracles (see section 2.5.1). We measure the complexity of our algorithms by the number of oracle calls and arithmetic computations.

Our first algorithm, based on Klartag’s M-Ellipsoid construction, yields a $\text{poly}(n)$ time randomized algorithm which succeeds with high probability.

Theorem 4.1.1 (M-Ellipsoid generator, informal). *There is a polynomial-time randomized algorithm that with high probability computes an M-Ellipsoid E of a given*

n -dimensional convex body K .¹

The next result is an algorithm which compute a near optimal covering of a convex body K by an a given ellipsoid E . In particular, it allows us to certify that a given ellipsoid E is an M-Ellipsoid of K .

Theorem 4.1.2 (Ellipsoid covering algorithm, informal). *Given an n dimensional convex body K and ellipsoid E , there is a deterministic $2^{O(n)}N(K, E)$ -time and polynomial space algorithm which outputs a covering of K by E of size at most $2^{O(n)}N(K, E)$.*²

We remark that the above algorithm outputs the elements of the covering one at a time without storing them; this is critical for ensuring the polynomial space complexity. Combining the previous two theorems (using a certain duality relation to bound $N(E, K)$), we get an expected $2^{O(n)}$ -time and polynomial space algorithm that is guaranteed to output an M-Ellipsoid and its implied covering for any given convex body K . We note that the computed covering of K by E will play a critical role in both the applications to volume estimation and lattices problems.

The second algorithm, based on Milman's M-Ellipsoid construction, yields a deterministic $2^{O(n)}$ time and $\text{poly}(n)$ space algorithm. Moreover, there is a $2^{\Omega(n)}$ lower bound for deterministic algorithms in the oracle model, so this is the best possible up to a constant in the exponent.

Theorem 4.1.3. *There is a deterministic $2^{O(n)}$ time and $\text{poly}(n)$ space algorithm that computes an M-Ellipsoid E of any given n dimensional convex body K .*

We remark that the above algorithm will not need to build coverings to certify the outputted M-Ellipsoid. However, as mentioned previously, the covering of K by E will be crucial for the intended applications.

¹ We thank Bo'az Klartag for suggesting to us that his M-Ellipsoid construction [78] can be made algorithmic.

² We thank Gideon Schechtman for suggesting the use of a parallelepiped tiling to compute the covering.

We now explore the consequences of the deterministic M-Ellipsoid construction to the problem of estimating the volume of a convex body. This is an ancient problem that has led to many insights and algorithmic techniques in high-dimensional geometry and probability theory. On the one hand, the problem can be solved for any convex body presented in the general membership oracle model in randomized polynomial time to arbitrary accuracy [45]. On the other hand, the following lower bound (improving on [48]) shows that deterministic algorithms cannot achieve such approximations.

Theorem 4.1.4. [51] *Suppose there is a deterministic algorithm that takes a convex body K as input and outputs $A(K), B(K)$ such that $A(K) \leq \text{vol}(K) \leq B(K)$ and makes at most n^a calls to the membership oracle for K . Then there is some convex body K for which*

$$\frac{B(K)}{A(K)} \leq \left(\frac{cn}{a \log n} \right)^{n/2}$$

where c is an absolute constant.

In particular, this implies that even achieving a $2^{O(n)}$ approximation requires $2^{\Omega(n)}$ oracle calls. Now the volume of an M-Ellipsoid E of K is clearly within a factor of $2^{O(n)}$ of the volume of K , thus Theorem 4.1.3 gives a $2^{O(n)}$ algorithm that achieves this approximation. And, as claimed, we have a lower bound of $2^{\Omega(n)}$ for computing an M-Ellipsoid deterministically. We state this corollary formally.

Theorem 4.1.5. *There is a deterministic $2^{O(n)}$ time and polynomial space algorithm that estimates the volume of any given n -dimensional convex body K to within a $2^{O(n)}$ multiplicative factor.*

A natural question is whether this can be generalized to a trade-off between approximation and complexity. Indeed the following result of Bárány and Füredi [52] gives a lower bound.

Theorem 4.1.6. [52] *For any $0 \leq \epsilon \leq 1$, any deterministic algorithm that estimates the volume of any input symmetric convex body to within a $(1 + \epsilon)^n$ factor given only a membership oracle to the body, must make at least $(1 + 1/\epsilon)^{\Omega(n)}$ queries to the membership oracle.*

Our next result answers this question in the affirmative. Here we apply the techniques from the deterministic M-Ellipsoid construction to give an algorithm that essentially matches the best possible complexity (up to a constant in the exponent) for centrally symmetric convex bodies.

Theorem 4.1.7. *For any $0 \leq \epsilon \leq 1$, there is a deterministic algorithm that given a symmetric convex body $K \subseteq \mathbb{R}^n$ computes a number $V \geq 0$, satisfying $V \leq \text{vol}_n(K) \leq (1 + \epsilon)^n V$, in $2^{O(n)}(1 + 1/\epsilon)^{O(n)}$ time and polynomial space.*

Our last result concerns finding a good “central” point for a convex body K . For symmetric convex bodies, the center of symmetry is the obvious center to pick. For asymmetric bodies however, the question becomes more subtle. For the lattice algorithms in the following chapters, a useful notion of good center, will be any point “approximately” maximizing the following ratio

$$\max_{\mathbf{x} \in K} \frac{\text{vol}_n((K - \mathbf{x}) \cap (\mathbf{x} - K))}{\text{vol}_n(K)}. \quad (4.1.1)$$

In convex geometry, the above ratio is as known Kovner-Besicovitch measure of symmetry for K . It is known that a uniform point $X \in K$ satisfies that the expected ratio is 2^{-n} . It was shown by Milman and Pajor [95] that this bounds holds in particular for the centroid of K (see Theorem 2.3.7).

We shall call any point $\mathbf{x} \in K$ for which the ratio (4.1.1) is $2^{-O(n)}$ an approximate center of mass. Using the properties of the centroid and our deterministic volume estimation algorithm, we give a Las Vegas algorithm for computing good approximate centers of mass.

Theorem 4.1.8. *For any convex body K , there is a randomized algorithm which runs in $2^{O(n)}$ time, using polynomial space and randomness that computes a center $\mathbf{x} \in K$ satisfying $\text{vol}_n((K - \mathbf{x}) \cap (\mathbf{x} - K)) \geq 5^{-n} \text{vol}_n(K)$.*

4.1.2 The M-Ellipsoid

An M-Ellipsoid of a convex body K is an ellipsoid E with the property that at most $2^{O(n)}$ translated copies of E are sufficient to cover all of K , and at most $2^{O(n)}$ copies of K are sufficient to cover E . The following theorem was first proved for symmetric bodies by Milman [92] and extended by Milman and Pajor [95] to the general case.

Theorem 4.1.9 ([95]). *There exists an absolute constant $C > 0$, such that for all $n \geq 1$ and any convex body $K \subseteq \mathbb{R}^n$, there exists an ellipsoid E satisfying*

$$N(K, E) \cdot N(E, K) \leq C^n. \quad (4.1.2)$$

Definition 4.1.10 (M-Ellipsoid). Let $K \subseteq \mathbb{R}^n$ be a convex body. If E is an ellipsoid satisfying Equation (4.1.2) (for some particular fixed C) with respect to K , then we say that E is an M-Ellipsoid of K .

There are many equivalent ways of understanding the M-Ellipsoid; here we list a few (proofs of many of these equivalences can be found in [95]).

Theorem 4.1.11. *Let $K \subseteq \mathbb{R}^n$ be a convex body with $b(K) = 0$ (centroid at the origin), and let $E \subseteq \mathbb{R}^n$ be an origin-centered ellipsoid. Then the following conditions are equivalent, where the absolute constant C may vary from line to line:*

- (1) $N(K, E) \cdot N(E, K) \leq C^n$.
- (2) $\text{vol}(K + E) \leq C^n \cdot \min\{\text{vol}(E), \text{vol}(K)\}$.
- (3) $\sup_{t \in \mathbb{R}^n} \text{vol}(K \cap (t + E)) \geq C^{-n} \cdot \max\{\text{vol}(E), \text{vol}(K)\}$.
- (4) E^* is an M-Ellipsoid of K^* .

From the above we see that the M-Ellipsoid is very robust object, and in particular is stable under polarity (assuming K is near-symmetric). We will use this fact (or a slight variant of it) in what follows, to help us certify a candidate M-Ellipsoid.

For specific examples, M-Ellipsoids for the ℓ_p balls are easy to describe. Using condition 3 of Theorem 4.1.11 above and standard volume estimates for ℓ_p balls, i.e., that $\text{vol}(B_p^n)^{1/n} = \Theta(n^{-1/p})$, we have the following:

Lemma 4.1.12. *Let B_p^n denote the n -dimensional ℓ_p ball. Then*

- *For $1 \leq p \leq 2$, $n^{\frac{1}{2}-\frac{1}{p}} \cdot B_2^n \subseteq B_p^n$ (the largest inscribed ball in B_p^n) is an M-Ellipsoid for B_p^n .*
- *For $p \geq 2$, $n^{\frac{1}{2}-\frac{1}{p}} \cdot B_2^n \supseteq B_p^n$ (the smallest containing ball of B_p^n) is an M-Ellipsoid for B_p^n .*

For general convex bodies, the first proof of existence for the M-Ellipsoid (see [92]) relies on a technique, developed by Milman, known as isomorphic symmetrization. This technique, which we describe and implement in section 4.4, slowly transforms any input body into an ellipsoid via a sequence of surgeries. Though the construction does not seem implementable in polynomial time, it can be used to yield a $2^{O(n)}$ time deterministic algorithm. As mentioned in the introduction, there is a $2^{\Omega(n)}$ lower bound for any such deterministic construction.

In section 4.3, we present a more direct construction of Klartag, which admits a randomized polynomial time algorithm. In contrast to the symmetrization approach, Klartag's approach is fundamentally randomized, and seems difficult to make deterministic.

To begin, in section 4.2, we give an efficient procedure to compute a covering of a convex body by an ellipsoid.

4.2 *Building a Covering*

In this section, we describe how to efficiently build a covering of a convex body $K \subseteq \mathbb{R}^n$ by an ellipsoid E . To construct such a covering, we first reduce the problem of covering K by E to the problem of computing a tiling of K by a maximum volume inscribed parallelepiped P of E . Though the such a tiling will clearly yield a covering of K by E (since $P \subseteq E$), it will not be optimal. However, we will see that the size of the tiling will be at worst a $2^{O(n)}$ factor larger than $N(K, E)$ (the size of the optimal covering of K by E), which will suffice for our purposes.

To understand the structure of such a tiling, we first note that the centers of the translates of P in the tiling of K correspond to points in a lattice. Due to this special lattice structure, we will be able to lazily enumerate the centers in the tiling (i.e. produce them one by one on demand) very efficiently using only polynomial space. Here we will rely on a space efficient graph enumeration technique known as reverse search, which was first developed by Avis and Fukuda [6] for enumerating the vertices of rational polyhedra.

Reverse Search for Enumeration: To be able to describe our parallelepiped tiling algorithm (which will be used to build ellipsoid coverings), we first introduce the Avis-Fukuda reverse search enumeration technique [6]³. This technique was developed by Avis and Fukuda to get an efficient algorithm for enumerating the vertices of a polyhedron or hyperplane arrangement.

In the reverse search setting, we start with a graph $G = (V, E)$ (generally with only an implicit description) of max degree Δ . To interact with G , we have access to an adjacency oracle Adj , which on input $\mathbf{v} \in V$ and k , $1 \leq k \leq \Delta$, either returns the index k neighbor of \mathbf{v} or returns NULL if \mathbf{v} has no neighbor at index k . Here we have

³We are indebted to Matthias Köppe for suggesting the use of reverse search to reduce the space complexity of our covering and enumeration algorithms.

the guarantee that the neighbors of \mathbf{v} returned from distinct indices are distinct, and that every neighbor of \mathbf{v} has an associated index.

Finally, we are given a local search function $f : V \rightarrow V$ of G . Let $S = \{\mathbf{v} \in V : f(\mathbf{v}) = \mathbf{v}\}$ denote the sink of nodes of f . By local search function of G , we mean that for any $\mathbf{v} \in V$ there exists a finite integer $k \geq 1$ such that $f^{(k)}(\mathbf{v}) = \mathbf{v}$ (where $f^{(k)}(\mathbf{v}) = f(f^{(k-1)}(\mathbf{v}))$ and $f^{(0)}(\mathbf{v}) = \mathbf{v}$), and that $\{f(\mathbf{v}), \mathbf{v}\} \subseteq E$ for all $\mathbf{v} \in V \setminus S$. Stated differently, the graph $T(f) = (V, E(f))$, where $E(f) = \{(\mathbf{v}, f(\mathbf{v})) : \mathbf{v} \in V \setminus S\}$, is directed subforest of G whose sinks correspond to S .

The main goal of reverse search is to “discover” the graph G , i.e. to output all the vertices of G in a time and space efficient manner. Given the adjacency oracle, we can certainly perform a breadth or depth search search of G to achieve this, however such an approach requires space proportional to the size of the graph (which could be huge) and hence is undesirable. To avoid this, Avis and Fukuda propose to use the directed tree structure $T(f)$ induced by f on G to perform the enumeration. When G is a rooted tree, a full traversal of the tree starting from the root, crossing each edge exactly twice, can be performed while only storing two nodes of the tree at any one time, i.e. the current visited node its immediate predecessor.

Using the local search and adjacency function, we will be able to traverse the tree $T(f)$ in the aforementioned way using only local information. In particular, given a vertex $\mathbf{v} \in V$, we will only need a way to list the parent and the child nodes of \mathbf{v} in $T(f)$. Here, the parent of \mathbf{v} in $T(f)$ is $f(\mathbf{v})$, and $\mathbf{w} = \text{Adj}(\mathbf{v}, k)$, $1 \leq k \leq \Delta$, is a child of \mathbf{v} in $T(f)$ iff $\mathbf{w} \neq \text{NULL}$ and $f(\mathbf{w}) = \mathbf{v}$. Now given a sink vertex $\mathbf{s} \in S$, reverse search will be able to enumerate the entire (weakly) connected component of $T(f)$ containing \mathbf{s} . The main caveat here is that the reverse search procedure requires the set of sink vertices S to be given explicitly to initiate the tree traversal. This must be achieved in an instance specific way.

We now provide an implementation of the Avis-Fukuda reverse search algorithm

(taken from [6]) and its associated guarantees.

Algorithm 4.1 Reverse-Search(Adj, Δ, S, f)

Input: Adjacency oracle Adj for graph $G = (V, E)$, max degree bound Δ , initial set of sinks $S \subseteq V$, and local search function f .

Output: Outputs V .

```

for each vertex  $s \in S$  do
   $v \leftarrow s; j \leftarrow 0$                                  $\triangleright j$ : neighbor counter
  repeat
    while  $j < \Delta$  do
      if  $j = 0$  then
        output  $v$                                            $\triangleright$  Output on first occurrence
         $j \leftarrow j + 1; \text{next} \leftarrow \text{Adj}[v, j]$ 
        if  $\text{next} \neq \text{NULL}$  and  $f(\text{next}) = v$  then
           $v \leftarrow \text{next}; j \leftarrow 0$                  $\triangleright$  Reverse traverse
        if  $v \neq s$  then
           $u \leftarrow v; v \leftarrow f(v); j \leftarrow 0$   $\triangleright$  Forward traverse
          repeat
             $j \leftarrow j + 1$ 
            until  $\text{Adj}[v, j] = u$                            $\triangleright$  Restore  $j$ 
          until  $v = s$  and  $j = \Delta$ 

```

Theorem 4.2.1 (Reverse Search [6]). *Given $G = (V, E)$, Adj , Δ , S and f as above, Algorithm 4.1 outputs the vertices of G (each vertex is outputted exactly once) using space proportional to storing 2 elements of V , using $O(\Delta|V|)$ queries to the Adj oracle and $O(|E|)$ queries to the local search function f . Furthermore, if the search is halted after having visited $N \leq |V|$ vertices of G , the number of calls to Adj and f up until termination is $O(\Delta N)$.*

We now present the Parallelepiped-Tiling algorithm which uses reverse search to construct a parallelepiped tiling of any convex body K in a time and space efficient manner.

Theorem 4.2.2. *Algorithm 4.2 Algorithm Parallelepiped-Tiling is correct, and runs in $4^n N(K, P)$ time using polynomial space.*

Proof. Since B is invertible, we note that $P = B[-1, 1]^n$ tiles \mathbb{R}^n with respect to the lattice $\mathcal{L} = 2B\mathbb{Z}^n$, i.e. $\mathcal{L} + P = \mathbb{R}^n$ and for distinct $\mathbf{x}, \mathbf{y} \in \mathcal{L}$, $\mathbf{x} + P \cap \mathbf{y} + P = \emptyset$.

Algorithm 4.2 Parallelepiped-Tiling(K, P, ϵ): Deterministic parallelepiped tiling of a convex body.

Input: A weak membership oracle O_K for an (\mathbf{a}_0, r, R) -centered convex body K , a parallelepiped $P = B[-1, 1]^n$, $B \in \mathbb{R}^{n \times n}$ invertible, and tolerance $0 < \epsilon \leq 1$.

Output: Outputs tiling Λ of K by P with respect to $\mathcal{L} = 2B\mathbb{Z}^n$, satisfying $\Lambda \subseteq K + (1 + \epsilon)P$.

- 1: Build intersection oracle INT, such that $\text{INT}(\mathbf{x}, \epsilon) = 1$ if $K \cap (P + \mathbf{x}) \neq \emptyset$ and $\text{INT}(\mathbf{x}, \epsilon) = 0$ if $K \cap ((1 + \epsilon)P + \mathbf{x}) = \emptyset$.
- 2: Let $\mathcal{L} = 2B\mathbb{Z}^n$, and let $\{\mathbf{u}_1, \dots, \mathbf{u}_{2n}\} = \{\pm 2B\mathbf{e}_1, \dots, \pm 2B\mathbf{e}_n\}$.
- 3: Let G be the graph with vertex set $V = \{\mathbf{x} \in \mathbf{a}_0 + \mathcal{L} : \text{INT}(\mathbf{x}, \epsilon) = 1\}$, and edge set $E = \{(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in V, \mathbf{x} - \mathbf{y} \in \{\mathbf{u}_1, \dots, \mathbf{u}_{2n}\}\}$.
- 4: Build adjacency oracle Adj for G , where for $\mathbf{y} \in V$ and $i \in [2n]$, $\text{Adj}[\mathbf{y}, i]$ returns $\mathbf{y} + \mathbf{u}_i$ if $\text{INT}(\mathbf{y} + \mathbf{u}_i, \epsilon) = 1$ and NULL otherwise.
- 5: Build local search function f for G , where for $\mathbf{y} \in V$, $f(\mathbf{y})$ either returns $\text{Adj}[\mathbf{y}, i]$ for the minimum $i \in [2n]$ satisfying $\text{Adj}[\mathbf{y}, i] \neq \text{NULL}$ and

$$\left\| \frac{1}{2}B^{-1}(\text{Adj}[\mathbf{y}, i] - \mathbf{a}_0) \right\|_1 = \left\| \frac{1}{2}B^{-1}(\mathbf{y} - \mathbf{a}_0) \right\|_1 - 1,$$

or returns \mathbf{y} if no such i exists.

- 6: **return** Reverse-Search(Adj, $2n, \mathbf{a}_0, f$)
-

We now wish to tile K with copies of P . We examine the set of centers of tiles intersecting K , i.e.

$$H = \{\mathbf{x} \in \mathbf{a}_0 + \mathcal{L} : (\mathbf{x} + P) \cap K \neq \emptyset\} = (\mathbf{a}_0 + \mathcal{L}) \cap (K + P)$$

We note that since K is (\mathbf{a}_0, r, R) centered, $\mathbf{a}_0 \in H$. Since $P + \mathcal{L} = \mathbb{R}^n$, it is easy to see that $K \subseteq H + P$. To successfully output the centers in the tiling H , we shall need to decide for $\mathbf{x} \in \mathbf{a}_0 + \mathcal{L}$, whether $\mathbf{x} + P \cap K \neq \emptyset$. For simplicity, in the rest of the proof, we replace $P = B[-1, 1]^n$ by $\bar{P} = B[-1, 1]^n$ (the closure of P). Clearly the set of intersecting tiles is only larger in this way, and hence the output will still enable us to compute a tiling for the original P .

Distinguishing Intersecting Tiles. Since we only have a weak membership oracle for K , we will only be able to decide whether $\mathbf{x} + P$ approximately intersects K . To formalize this, we build an weak intersection oracle INT which queried on $\mathbf{x} \in \mathbb{R}^n$,

$\epsilon > 0$ satisfies

$$\text{INT}(\mathbf{x}, \epsilon) = \begin{cases} 0 & : \quad \mathbf{x} + P \cap K = \emptyset \\ 1 & : \quad \mathbf{x} + (1 + \epsilon)P \cap K \neq \emptyset \end{cases}.$$

Using this oracle we will be able to overestimate H , and compute a set $S \subseteq \mathbf{a}_0 + \mathcal{L}$ such that

$$H \subseteq S \subseteq \{\mathbf{x} \in \mathbf{a}_0 + \mathcal{L} : \mathbf{x} + (1 + \epsilon)P \cap K \neq \emptyset\}$$

which will suffice for our purposes. Now to build INT, we first remark that for $\mathbf{x} \in \mathbb{R}^n$, $t \geq 0$

$$\mathbf{x} + tP \cap K \neq \emptyset \Leftrightarrow \inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P \leq t \Leftrightarrow \inf_{\mathbf{y} \in K} \|B^{-1}(\mathbf{y} - \mathbf{x})\|_\infty \leq t$$

Hence deciding the minimum scaling t of P for which $\mathbf{x} + tP \cap K \neq \emptyset$ is equivalent to solving a simple convex program. The above convex program is of the form described in Theorem 2.5.9, hence for $\epsilon > 0$, and $\mathbf{x} \in \mathbb{Q}^n$, we may compute a number $\omega \geq 0$ such that

$$|\omega - \inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P| \leq \epsilon \tag{4.2.1}$$

in polynomial time. We now build INT. On query $\mathbf{x} \in \mathbb{Q}^n$, $\epsilon > 0$, we do the following:

- (1) Compute $\omega \geq 0$ satisfying $|\omega - \inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P| \leq \frac{\epsilon}{2}$.
- (2) If $\omega \leq 1 + \frac{\epsilon}{2}$ return 1, otherwise return 0.

From (4.2.1) the above procedure clearly runs in polytime. To prove correctness, we must show that $\text{INT}(\mathbf{x}, \epsilon) = 1$ if $\mathbf{x} + P \cap K \neq \emptyset$ and $\text{INT}(\mathbf{x}, \epsilon) = 0$ if $\mathbf{x} + (1 + \epsilon)P \cap K = \emptyset$. If $(\mathbf{x} + P) \cap K \neq \emptyset$, we note that $\inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P \leq 1$, hence by the guarantee on ω we have that

$$\omega \leq \inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P + \frac{\epsilon}{2} \leq 1 + \frac{\epsilon}{2},$$

and so we correctly classify \mathbf{x} . If $\mathbf{x} + (1 + \epsilon)P \cap K = \emptyset$, then $\inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P > 1 + \epsilon$ and so

$$\omega \geq \inf_{\mathbf{y} \in K} \|\mathbf{y} - \mathbf{x}\|_P - \frac{\epsilon}{2} > 1 + \frac{\epsilon}{2}$$

as needed.

Implementing Reverse Search to Compute the Tiling. Let

$$H_\epsilon = \{\mathbf{x} \in \mathbf{a}_0 + \mathcal{L} : \text{INT}(\mathbf{x}, \epsilon) = 1\}.$$

By construction of the intersection oracle INT, we have that $H \subseteq H_\epsilon \subseteq (a_0 + \mathcal{L}) \cap (K + (1 + \epsilon)P)$. Hence the elements of H_ϵ contain the desired tiling H with a “small” number of extraneous tiles.

To compute the tiling, we will run a tailored reverse search on the graph $G = (V, E)$ where $V = H_\epsilon$ and $E = \{\{\mathbf{x}, \mathbf{y}\} : \mathbf{x}, \mathbf{y} \in V, \mathbf{x} - \mathbf{y} \in 2B\{\pm\mathbf{e}_1, \dots, \pm\mathbf{e}_n\}\}$. Note that the G has max degree $\Delta = 2n$. To implement the adjacency oracle Adj, we first let $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{2n}$ denote an ordering of the vectors $\pm 2B\mathbf{e}_1, \dots, \pm 2B\mathbf{e}_n$. For a vertex $\mathbf{v} \in H_\epsilon$, and integer k , $1 \leq k \leq 2n$, $\text{Adj}(\mathbf{v}, k)$ returns $\mathbf{v} + \mathbf{u}_k$ if $\text{INT}(\mathbf{v} + \mathbf{u}_k, \frac{1}{n}) = 1$ (i.e. checks whether $\mathbf{v} + \mathbf{u}_k \in H_\epsilon$) and NULL otherwise. We note that Adj is correct and runs in polynomial time.

Next we implement a local search function f satisfying the following properties: (1) $\mathbf{a}_0 \in H_\epsilon$ is a sink of f , (2) the connected component of \mathbf{a}_0 in $T(f)$ (the directed subforest induced by f) contains H , and (3) f runs in polynomial time. We note that if we build f satisfying properties (1) and (2), then running reverse search from the root \mathbf{a}_0 gives us a superset of the desired tiling H . On input $v \in H_\epsilon$, the local search f does the following. Compute the minimum $i \in [2n]$, such that $\mathbf{v} + \mathbf{u}_i \in H_\epsilon$ and

$$\left\| \frac{1}{2}B^{-1}(\mathbf{v} + \mathbf{u}_i - \mathbf{a}_0) \right\|_1 = \left\| \frac{1}{2}B^{-1}(\mathbf{v} - \mathbf{a}_0) \right\|_1 - 1.$$

If a minimizing index i satisfying the above is found, return $\mathbf{v} + \mathbf{u}_i$; else, return \mathbf{v} . Here we see that f attempts to find the lowest indexed neighbor which gets closer to the “root” \mathbf{a}_0 (under a certain ℓ_1 distance). To output the desired tiling, we run the reverse search algorithm on the graph G , Δ , adjacency oracle Adj, local search function f , and initial set $S = \{\mathbf{a}_0\}$.

The correctness of the algorithm at this point depends solely on the stated properties of the local search function f . These are proved in the following claim.

Claim: f satisfies properties (1),(2) and (3).

Proof. Without loss of generality (after an appropriate affine transformation), we may assume that $B = \frac{1}{2}I_n$ (I_n is $n \times n$ identity) and that $\mathbf{a}_0 = \mathbf{0}$ (and hence $\mathbf{0} \in K$). In this case, we see that $\mathcal{L} = 2B\mathbb{Z}^n = \mathbb{Z}^n$, $P = [-1/2, 1/2]^n$, $H = (K + P) \cap \mathbb{Z}^n$.

For (1), note that any neighbor \mathbf{w} of $\mathbf{0}$ in G (i.e. one of $\pm\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$) is non-zero and integral, and hence satisfies $\|\frac{1}{2}B^{-1}\mathbf{w}\|_1 = \|\mathbf{w}\|_1 \geq 1$. Therefore, by construction of f , we have that $f(\mathbf{0}) = \mathbf{0}$ as needed.

For (2), we first show that every $\mathbf{v} \in H$, $\mathbf{v} \neq \mathbf{0}$, has a neighbor \mathbf{w} in G such that $\|\mathbf{w}\|_1 = \|\mathbf{v}\|_1 - 1$. Since $\mathbf{v} \in H$, we have that $(\mathbf{v} + P) \cap K \neq \emptyset$. Pick $\mathbf{y} \in (\mathbf{v} + P) \cap K$. Note that

$$\mathbf{v} + P = \mathbf{v} + [-1/2, 1/2]^n = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{u}_i^t(\mathbf{x} - \mathbf{v}) \leq \frac{1}{2}, i \in [2n]\}$$

where $\{\mathbf{u}_1, \dots, \mathbf{u}_{2n}\} = \{\pm\mathbf{e}_1, \dots, \pm\mathbf{e}_n\}$. Examine the line $(1 - \alpha)\mathbf{y}$ for $\alpha \in [0, 1]$. Since $\mathbf{0} \notin \mathbf{v} + P$ and $\mathbf{v} + P$ is a closed and convex, there exists a minimum $\alpha \in [0, 1)$ such that $\bar{\mathbf{y}} = (1 - \alpha)\mathbf{y} \in \mathbf{v} + P$ and $\mathbf{y} - \delta\mathbf{y} \notin \mathbf{v} + P$ for all $\delta > 0$. Since $\mathbf{0}, \mathbf{y} \in K$, we see that $\bar{\mathbf{y}} \in K$. Note that $\bar{\mathbf{y}}$ is on the boundary of $\mathbf{v} + P$, and so the set $I = \{i \in [2n] : \mathbf{u}_i^t(\bar{\mathbf{y}} - \mathbf{v}) = \frac{1}{2}\}$ (indices of the tight constraints for $\bar{\mathbf{y}}$) is non-empty. I claim that there exists $i \in I$, such that $\mathbf{u}_i^t\bar{\mathbf{y}} < 0$. Assume not, then for $\delta > 0$ and $i \in I$, $\mathbf{u}_i^t(\bar{\mathbf{y}} - \delta\mathbf{y} - \mathbf{v}) = \frac{1}{2} - \delta\mathbf{u}_i^t\bar{\mathbf{y}} \leq \frac{1}{2}$. Since the constraints indexed by $i \in [2n] \setminus I$ are not tight at $\bar{\mathbf{y}}$, there exists $\delta > 0$ (small enough) such that $\bar{\mathbf{y}} - \delta\mathbf{y} \in \mathbf{v} + P$, a contradiction to our assumption on $\bar{\mathbf{y}}$.

Take $i \in I$ such that $\mathbf{u}_i^t\bar{\mathbf{y}} < 0$. I claim that $\mathbf{y} + \mathbf{u}_i \in H$, and that $\|\mathbf{y} + \mathbf{u}_i\|_1 = \|\mathbf{y}\|_1 - 1$. First, we show that $\bar{\mathbf{y}} \in \mathbf{u}_i + \mathbf{v} + P = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{u}_j^t(x - \mathbf{u}_i - \mathbf{v}) \leq \frac{1}{2}, j \in [2n]\}$. Since the \mathbf{u}_j 's are sign flips of the standard basis vectors, we have that $\mathbf{u}_j^t\mathbf{u}_i = 0$ unless $\mathbf{u}_j = \pm\mathbf{u}_i$. For $j \in [2n]$ satisfying $\mathbf{u}_j^t\mathbf{u}_i = 0$, note that $\mathbf{u}_j^t(\bar{\mathbf{y}} - \mathbf{u}_i - \mathbf{v}) = \mathbf{u}_j^t(\bar{\mathbf{y}} - \mathbf{v}) \leq \frac{1}{2}$ since $\bar{\mathbf{y}} \in \mathbf{v} + P$. Now assume $\mathbf{u}_j = \mathbf{u}_i$, then $\mathbf{u}_j^t(\bar{\mathbf{y}} - \mathbf{u}_i - \mathbf{v}) = \frac{1}{2} - \|\mathbf{u}_i\|^2 = -\frac{1}{2}$. Lastly, if $\mathbf{u}_j = -\mathbf{u}_i$, then $-\mathbf{u}_j^t(\bar{\mathbf{y}} - \mathbf{u}_i - \mathbf{v}) = \frac{1}{2}$. Hence $\bar{\mathbf{y}} \in \mathbf{v} + \mathbf{u}_i + P$ as needed. Since

$\bar{\mathbf{y}} \in K$ we have that $\mathbf{v} + \mathbf{u}_i \in K + P$, and hence $\mathbf{v} + \mathbf{u}_i \in H$. Lastly, we show that $\|\mathbf{v} + \mathbf{u}_i\|_1 = \|\mathbf{v}\|_1 - 1$. Take $j \in [n]$, such that $\mathbf{u}_i = \pm \mathbf{e}_j$. Here we have that

$$\begin{aligned} \|\mathbf{v} + \mathbf{u}_i\|_1 &= \sum_{k \in [n]} |\mathbf{e}_k^t(\mathbf{v} + \mathbf{u}_i)| = |\mathbf{e}_j^t(\mathbf{v} + \mathbf{u}_i)| + \sum_{k \in [n], k \neq j} |\mathbf{e}_k^t \mathbf{v}| \\ &= |\mathbf{e}_j^t(\mathbf{v} + \mathbf{u}_i)| - |\mathbf{e}_j^t \mathbf{v}| + \|\mathbf{v}\|_1 \end{aligned}$$

Let $a = \mathbf{e}_j^t \mathbf{v}$ and $b = \mathbf{e}_j^t \mathbf{u}_i$. Since $\mathbf{v} \in \mathbb{Z}^n$ and $\mathbf{u}_i = \pm \mathbf{e}_j$, we have that $a \in \mathbb{Z}$ and $b \in \{-1, 1\}$. I claim that $ab < 0$. To see this, first note that $\mathbf{u}_i^t(\bar{\mathbf{y}} - \mathbf{v}) = \frac{1}{2}$ and hence $\mathbf{u}_i^t(\bar{\mathbf{y}}) - \frac{1}{2} = \mathbf{u}_i^t \mathbf{v} = ab$. Since $\mathbf{y}^t \mathbf{u}_i < 0$ and $\bar{\mathbf{y}} = (1 - \delta)\mathbf{y}$, $\delta \in [0, 1)$, we have that $\mathbf{u}_i^t \bar{\mathbf{y}} < 0$. Hence $ab < -\frac{1}{2} < 0$. Combining $ab < 0$, $a \in \mathbb{Z}$ and $b \in \{-1, 1\}$, yields that $|\mathbf{e}_j(\mathbf{v} + \mathbf{u}_i)| = |a + b| = |a| - 1$. Therefore $\|\mathbf{v} + \mathbf{u}_i\|_1 = \|\mathbf{v}\|_1 - 1$ as needed.

To prove (2), we note that for $\mathbf{v} \in \mathbb{Z}^n$, $\|\mathbf{v}\|_1 = \sum_{i \in [n]} |\mathbf{e}^t \mathbf{v}| \in \mathbb{Z}$. From the above argument, for $\mathbf{v} \in H = (K + P) \cap \mathbb{Z}^n$, $\mathbf{v} \neq \mathbf{0}$, we have that $f(\mathbf{v}) \in H$ and $\|f(\mathbf{v})\|_1 = \|\mathbf{v}\|_1 - 1$. Therefore letting $t = \|\mathbf{v}\|_1 \in \mathbb{Z}$, we have that $\|f^{(t)}(\mathbf{v})\|_1 = \|\mathbf{v}\|_1 - t = 0$, and hence $f^{(t)}(\mathbf{v}) = \mathbf{0}$. Therefore the connected component of $\mathbf{0}$ in $T(f)$ contains H as needed.

Lastly, on input \mathbf{v} , f calls the oracle INT at most $2n$ times, and then performs basic matrix operations. Therefore f runs in polynomial time and satisfies (3) as needed. \square

Runtime: To bound the running time, we need only bound the time needed to run the reverse search. By Theorem 4.2.1, we have that the search uses $O(\Delta|V|) = O(2n|H_\epsilon|)$ calls to the adjacency oracle Adj and $O(|E|) = O(\Delta|V|) = O(2n|H_\epsilon|)$ queries to the local search function f . Now we note that

$$|H_\epsilon| = \frac{\text{vol}_n(H_\epsilon + P)}{\text{vol}_n(P)} \leq \frac{\text{vol}_n(K + (1 + \epsilon)P + P)}{\text{vol}_n(P)} = \frac{\text{vol}_n(K + (2 + \epsilon)P)}{\text{vol}_n(P)}$$

Now let $T \subseteq \mathbb{R}^n$ denote a covering of K by P (i.e. $K \subseteq T + P$), satisfying $|T| = N(K, P)$. Now note that

$$\frac{\text{vol}_n(K + (2 + \epsilon)P)}{\text{vol}_n(P)} \leq \frac{\text{vol}_n(T + (3 + \epsilon)P)}{\text{vol}_n(P)} \leq |T| \frac{\text{vol}_n((3 + \epsilon)P)}{\text{vol}_n(P)} = N(K, P)(3 + \epsilon)^n$$

Since both queries to f and Adj take polynomial time, the total running time is $4^n N(K, P) \text{poly}(\cdot)$ ($\epsilon \leq 1$). Lastly, since reverse search uses space polylogarithmic in $|H_\epsilon|$ (i.e. storing two nodes of the graph), the total space usage is polynomial. \square

We now give the straightforward reduction from ellipsoid covering to parallelepiped covering.

Algorithm 4.3 Ellipsoid-Cover(K, E): Deterministic construction of an ellipsoid covering of a convex body.

Input: A weak membership oracle O_K for an (a_0, r, R) -centered convex body K , an ellipsoid $E = E(A)$, $A \succ 0$.

Output: Outputs a covering of K by E of size at most $(3\sqrt{\frac{\pi e}{2}}(1 + o(1)))^n N(K, E)$.

- 1: Compute $B = A^{-\frac{1}{2}}$ and let $P = \frac{1}{\sqrt{n}}B[-1, 1]^n$ (a maximum-volume inscribed parallelepiped of E).
 - 2: **return** Parallelepiped-Tiling($K, P, \frac{1}{n}$)
-

Theorem 4.2.3. *Algorithm Ellipsoid-Cover is correct and runs in*

$(3\sqrt{\frac{\pi e}{2}}(1 + o(1)))^n N(K, E) \text{poly}(\cdot)$ time using polynomial space.

Proof.

Correctness: Given the correctness of algorithm Parallelepiped-Tiling, to check that the output is indeed a covering of K by E , we need only check that $P \subseteq E$.

To see this, we first note that $E = A^{-\frac{1}{2}}B_2^n$. Next by the containment $\frac{1}{\sqrt{n}}[-1, 1]^n = \frac{1}{\sqrt{n}}B_\infty^n \subseteq B_2^n$, we have that $P \subseteq \frac{1}{\sqrt{n}}A^{-\frac{1}{2}}B_\infty^n \subseteq A^{-\frac{1}{2}}B_2^n = E$, as needed.

Now we need to show that the size of the outputted covering is bounded by $N(K, E)$ in the appropriate way. To begin, we note that

$$\frac{\text{vol}_n(E)}{\text{vol}_n(P)} = \frac{\text{vol}_n(A^{-\frac{1}{2}}B_2^n)}{\text{vol}_n(\frac{1}{\sqrt{n}}A^{-\frac{1}{2}}B_\infty^n)} \leq \frac{\text{vol}_n(B_2^n)}{\text{vol}_n(\frac{1}{\sqrt{n}}B_\infty^n)} = \left(\sqrt{\frac{\pi e}{2}}(1 + o(1)) \right)^n$$

By the guarantees on the algorithm Parallelepiped-Tiling, on inputs $K, P, \frac{1}{n}$, it returns a tiling H of K by P satisfying $H \subseteq K + (1 + \frac{1}{n})P$. From the analysis of the

Parallelepiped-Tiling algorithm, we have that

$$\begin{aligned} |H| &\leq \frac{\text{vol}_n(K + (2 + \frac{1}{n})P)}{\text{vol}_n(P)} \leq \left(\sqrt{\frac{\pi e}{2}}(1 + o(1)) \right)^n \frac{\text{vol}_n(K + (2 + \frac{1}{n})P)}{\text{vol}_n(E)} \\ &\leq \left(\sqrt{\frac{\pi e}{2}}(1 + o(1)) \right)^n \frac{\text{vol}_n(K + (2 + \frac{1}{n})E)}{\text{vol}_n(E)} \end{aligned}$$

Now, let T denote a covering of K by E satisfying $|T| = N(K, E)$. Then

$$\frac{\text{vol}_n(K + (2 + \frac{1}{n})E)}{\text{vol}_n(E)} \leq \frac{\text{vol}_n(T + (3 + \frac{1}{n})E)}{\text{vol}_n(E)} \leq N(K, E) \left(3 + \frac{1}{n} \right)^n \leq N(K, E) 3^n e$$

This gives the final bound $|H| \leq \left(3\sqrt{\frac{\pi e}{2}}(1 + o(1)) \right)^n N(K, E)$ as needed.

Runtime: Since $A \succ 0$ we have $A^{-\frac{1}{2}}$ is well defined and can be computed in polynomial time. By the analysis of algorithm parallelepiped tiling, we know that the tiling algorithm runs in $|H| \text{poly}(\cdot)$ using polynomial space. Using the the above bound on $|H|$ yields the result. □

4.3 Klartag's Construction

Our first algorithm for generating a candidate M-Ellipsoid is based on a constructive proof of Theorem 4.1.9 by Klartag [78], who suggested to us the idea of using these techniques to build an M-Ellipsoid algorithmically.

Let $K \subseteq \mathbb{R}^n$ denote a convex body. To understand Klartag's construction, we begin with the assertion that under the *slicing conjecture* (also known as the *hyperplane conjecture*), a \sqrt{n} scaling of K 's inertial ellipsoid is an M-Ellipsoid — indeed, this is an equivalent form of the slicing conjecture (see [96] for a proof). Since the validity of the slicing conjecture is unknown, Klartag shows that a random perturbation K' of K , which remains close to K and has bounded isotropic constant. We reproduce the main theorem of [78] below:

Theorem 4.3.1 ([78]). *Let $K \subseteq \mathbb{R}^n$ be a convex body. Then for every real $\epsilon \in (0, 1)$, there exists a convex body $K' \subseteq \mathbb{R}^n$ such that*

- (1) $d(K, K') = \inf \{b/a : \exists \mathbf{x}, \mathbf{y} \in \mathbb{R}^n \text{ s.t. } a(K' - \mathbf{x}) \subseteq K - \mathbf{y} \subseteq b(K' - \mathbf{x})\} \leq 1 + \epsilon.$
- (2) $L_{K'} \leq c/\sqrt{\epsilon}.$

where $c > 0$ is an absolute constant and $L_{K'}$ is the isotropic constant of K' .

The relationship between the above theorem and the existence of the M-Ellipsoid is straightforward. First, from the closeness of K and K' it follows that an M-Ellipsoid for K' is an M-Ellipsoid for K . Lastly, from the bound on $L_{K'}$, a \sqrt{n} scaling of the inertial ellipsoid of K' is an M-Ellipsoid for K' .

Here we will not need to construct K' itself, but only an ellipsoid very close to its inertial ellipsoid (which as just mentioned is an M-Ellipsoid for K). The body K' is derived from a certain family of reweighted densities over K . These densities are given by exponential reweightings of the uniform density along some vector $\mathbf{s} \in \mathbb{R}^n$, i.e., $f_{\mathbf{s}}(\mathbf{x}) = e^{\langle \mathbf{s}, \mathbf{x} \rangle}$ for $\mathbf{x} \in K$ (and 0 otherwise). For \mathbf{s} chosen uniformly from $n \cdot \text{conv}\{K - \mathbf{b}(K), \mathbf{b}(K) - K\}^*$, the reweighting $f_{\mathbf{s}}$ has two important properties: (i) it is not too highly biased away from uniform over K , and (ii) it has bounded isotropic constant (independent of n) with very high probability. Let E be the inertial ellipsoid of $f_{\mathbf{s}}$ (or any reasonably good approximation to it), which can be found by sampling from $f_{\mathbf{s}}$. The first property of $f_{\mathbf{s}}$ allows us to prove that E can be covered by $2^{O(n)}$ copies of K , while the second property lets us cover K by $2^{O(n)}$ copies of E (see Lemma 4.3.8).

To make everything work algorithmically, we need robust versions of Klartag's main lemmas, since we will only be able to approximate the centroid of K , sample \mathbf{s} from near uniform distribution, and estimate the covariance matrix of $f_{\mathbf{s}}$.

Algorithm 4.4 makes the above description more formal. Note that given an oracle for a convex body, an oracle for the polar body can be constructed in polynomial time [56]. Sampling, both from the uniform and exponentially reweighted distributions, can be done in polynomial time using the random walk algorithm of [87, 86].

Theorem 4.3.8 together with Theorem 4.3.9 implies that the algorithm's output is indeed an M-Ellipsoid with good probability.

Algorithm 4.4 M-Gen: Randomized generation of a candidate M-Ellipsoid.

Input: A weak membership oracle O_K for a (\mathbf{a}_0, r, R) -centered convex body K .

Output: With probability $1 - o(1)$, an M-Ellipsoid E of K , or FAIL.

- 1: Estimate the centroid $\mathbf{b} = \mathbf{b}(K)$ using Algorithm Estimate-Centroid.
If Estimate-Centroid fails, return FAIL.
 - 2: Construct a membership oracle for $n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^*$.
 - 3: Sample a random vector \mathbf{s} from $n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^*$.
 - 4: Estimate the covariance matrix A of the density proportional to $e^{(\mathbf{s}, \mathbf{x})}$ on K .
 - 5: Output the ellipsoid $\sqrt{n}E(A^{-1}) = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^t A^{-1} \mathbf{x} \leq n\}$.
-

Theorem 4.3.2 (Correctness of M-Gen). *For large enough n , Algorithm 4.4 (M-Gen) outputs an ellipsoid E satisfying*

$$N(E, K) \leq (25e)^n \quad \text{and} \quad N(K, E) \leq (13e)^n \quad (4.3.1)$$

with probability at least $1 - \frac{1}{n}$ in polynomial time.

Proof of Theorem 4.3.2 (Correctness of M-Gen). The proof has two parts, first estimating the centroid of K and using it to build a membership oracle for the polar, and finally using this oracle to sample and estimate an appropriate inertial ellipsoid.

Estimating The Centroid: In the first step, we call algorithm Estimate-Centroid (Lemma 4.3.5) on K with failure probability guarantee $\frac{1}{4n}$. If Estimate-Centroid returns FAIL, we return FAIL. Else, Estimate-Centroid returns an estimate \mathbf{b} of $\mathbf{b}(K)$ with the guarantee

$$\frac{r}{2(n+1)\sqrt{n}}B_2^n \subseteq K - \mathbf{b} \subseteq 2RB_2^n$$

Furthermore, with probability at least $1 - \frac{1}{4n}$, we have that

$$\mathbf{b} - \mathbf{b}(K) \in \frac{1}{n+1}E_K \quad (4.3.2)$$

Building a membership oracle for the polar: From the guarantees on the algorithm Estimate-Centroid, we know that $K - \mathbf{b}$ is $(\mathbf{0}, \frac{r}{2(n+1)\sqrt{n}}, 2R)$ centered. We note these guarantees are polynomial in the input. Using these guarantees, we will build a polynomial time weak membership oracle for $S = n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^*$. We note that

$$\mathbf{v} \in n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^* \Leftrightarrow \max \left\{ \sup_{\mathbf{x} \in K - \mathbf{b}} \langle \mathbf{v}, \mathbf{x} \rangle, \sup_{\mathbf{x} \in K - \mathbf{b}} \langle -\mathbf{v}, \mathbf{x} \rangle \right\} \leq n$$

Given the guarantees on $K - \mathbf{b}$, we have that

$$\frac{n}{2R} B_2^n \subseteq n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^* \subseteq \frac{2n(n+1)\sqrt{n}}{r} B_2^n$$

Using the above characterization, and the sandwiching bounds, we will be able to build a weak membership oracle for S by approximately maximizing \mathbf{v} and $-\mathbf{v}$ over $K - \mathbf{b}$. For $\mathbf{v} \in \mathbb{R}^n$, $\epsilon > 0$, the weak membership oracle $O_S(\mathbf{v}, \epsilon)$ performs the following. If $\|\mathbf{v}\|_2 \leq \frac{n}{2R}$, return 1, if $\|\mathbf{v}\|_2 > \frac{2n(n+1)\sqrt{n}}{r}$ return 0, otherwise continue. Using the ellipsoid algorithm for convex optimization (Theorem 2.5.9), compute a number $\omega \geq 0$ satisfying

$$\omega - \frac{\epsilon r}{2(n+1)\sqrt{n}} \leq \max \left\{ \sup_{\mathbf{x} \in K - \mathbf{b}} \langle \mathbf{v}, \mathbf{x} \rangle, \sup_{\mathbf{x} \in K - \mathbf{b}} \langle -\mathbf{v}, \mathbf{x} \rangle \right\} \leq \omega.$$

If $\omega \leq n$, return 1, and otherwise return 0. Since ω can be computed in polynomial time, we see that O_S executes in polynomial time.

Claim: O_S yields a weak membership oracle for S .

Proof. We first show that if $\mathbf{v} \in S^{-\epsilon}$ then $O(\mathbf{v}, \epsilon) = 1$. Given that we accept if $\|\mathbf{v}\|_2 \leq \frac{n}{2R}$, we may assume that $\|\mathbf{v}\|_2 > \frac{n}{2R} > 0$. Since $\mathbf{v} \in S^{-\epsilon}$, we have that $\mathbf{v} + \epsilon \mathbf{v} / \|\mathbf{v}\|_2 \in S$. We recall the support function of $K - \mathbf{b}$, $h_{K - \mathbf{b}}(\mathbf{v}) = \sup_{\mathbf{x} \in K - \mathbf{b}} \langle \mathbf{v}, \mathbf{x} \rangle$.

Since $\epsilon \mathbf{v} / \|\mathbf{v}\|_2$ corresponds to a positive scaling of \mathbf{v} , we have that

$$\begin{aligned} n &\geq \max\{h_{K-\mathbf{b}}(\mathbf{v} + \epsilon \mathbf{v} / \|\mathbf{v}\|_2), h_{K-\mathbf{b}}(-\mathbf{v} - \epsilon \mathbf{v} / \|\mathbf{v}\|_2)\} \\ &= \max\{h_{K-\mathbf{b}}(\mathbf{v}) + \epsilon h_{K-\mathbf{b}}(\mathbf{v} / \|\mathbf{v}\|_2), h_{K-\mathbf{b}}(-\mathbf{v}) + \epsilon h_{K-\mathbf{b}}(-\mathbf{v} / \|\mathbf{v}\|_2)\} \\ &\geq \max\{h_{K-\mathbf{b}}(\mathbf{v}), h_{K-\mathbf{b}}(-\mathbf{v})\} + \epsilon \min\{h_{K-\mathbf{b}}(\mathbf{v} / \|\mathbf{v}\|_2), h_{K-\mathbf{b}}(-\mathbf{v} / \|\mathbf{v}\|_2)\} \end{aligned}$$

Since $\frac{r}{2(n+1)\sqrt{n}} B_2^n \subseteq K - \mathbf{b}$, we see that

$$\epsilon \min\{h_{K-\mathbf{b}}(\mathbf{v} / \|\mathbf{v}\|_2), h_{K-\mathbf{b}}(-\mathbf{v} / \|\mathbf{v}\|_2)\} \geq \epsilon \frac{r}{2(n+1)\sqrt{n}}$$

Combining the above inequalities, we get

$$n - \epsilon \frac{r}{2(n+1)\sqrt{n}} \geq \max\{h_{K-\mathbf{b}}(\mathbf{v}), h_{K-\mathbf{b}}(-\mathbf{v})\}$$

From here, we see that the number $\omega \geq 0$ computed by $O_S(\mathbf{v}, \epsilon)$ satisfies

$$\omega \leq \max\{h_{K-\mathbf{b}}(\mathbf{v}), h_{K-\mathbf{b}}(-\mathbf{v})\} + \epsilon \frac{r}{2(n+1)\sqrt{n}} \leq n$$

Hence $O_S(\mathbf{v}, \epsilon)$ returns 1 as needed.

Lastly if $\mathbf{v} \notin S^\epsilon$ we show that $O(\mathbf{v}, \epsilon) = 0$. Since $\mathbf{v} \notin S^\epsilon$, and S is closed, we have that $n < \max\{h_{K-\mathbf{b}}(\mathbf{v}), h_{K-\mathbf{b}}(-\mathbf{v})\}$. By the guarantee on ω , we have that $\omega \geq \max\{h_{K-\mathbf{b}}(\mathbf{v}), h_{K-\mathbf{b}}(-\mathbf{v})\} > n$. Therefore $O_S(\mathbf{v}, \epsilon) = 0$ as needed. \square

Building the M-Ellipsoid: Let π_S denote the uniform distribution on S . Equipped with a weak membership oracle for S , we may use the sampling algorithm of Theorem 2.5.11, to sample a point $Y \in S$ with distribution σ satisfying $d_{\text{TV}}(\sigma, \pi_S) \leq \frac{1}{4n}$ in polynomial time. Set $\mathbf{s} = Y$, where Y is the computed sample. We shall use \mathbf{s} to specify a reweighting of the uniform distribution on $K - \mathbf{b}$. Let $f_{\mathbf{s}}(\mathbf{x}) = e^{\langle \mathbf{s}, \mathbf{x} \rangle}$ for $\mathbf{x} \in K - \mathbf{b}$ and 0 otherwise. Using the algorithm described by Corollary 4.3.4, we may compute a matrix $A \in \mathbb{R}^{n \times n}$ satisfying

$$e^{-\frac{1}{n}} E_{f_{\mathbf{s}}} \subseteq E(A) \subseteq e^{\frac{1}{n}} E_{f_{\mathbf{s}}} \quad (4.3.3)$$

with probability $1 - \frac{1}{4n}$ in polynomial time. We return the ellipsoid $\sqrt{n}E(A)$ as our candidate M -ellipsoid for K .

Analysis: We now show that for n large enough, the ellipsoid returned by this algorithm satisfies with high probability the covering conditions

$$N(K, \sqrt{n}E(A)) \leq (13e)^n \quad \text{and} \quad N(\sqrt{n}E(A), K) \leq (25e)^n$$

First, we condition on the event (4.3.2), i.e. that we get a good estimate \mathbf{b} of $\mathbf{b}(K)$.

Next, we condition on the event (4.3.3), i.e. that we get a good estimate of E_{f_s} .

Hence at this point, our success probability is at least $1 - \frac{1}{2n}$.

Let $\eta > 0$ be a constant to be decided later. Let X be uniformly distributed on S , and let Y denote the approximately uniform sample the above algorithm computes on S , remembering that $S = n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^*$. Given the guarantee that $\mathbf{b}(K) - \mathbf{b} \in \frac{1}{n+1}E_K$, from Theorem 4.3.9 setting $\epsilon = 1$, for n large enough we have that

$$\mathbb{E}[L_{f_X}^{2n}] \leq \left((1 + o(1)) \sqrt{\frac{2}{\pi e}} \frac{e^\epsilon}{\sqrt{\epsilon}} \right)^{2n} \leq \left((1 + \eta) \sqrt{\frac{2e}{\pi}} \right)^{2n}$$

Using Markov's inequality, we see that

$$\Pr \left[L_{f_X} > (1 + \eta)^2 \sqrt{\frac{2e}{\pi}} \right] \leq \frac{\mathbb{E}[L_{f_X}^{2n}]}{\left((1 + \eta)^2 \sqrt{\frac{2e}{\pi}} \right)^{2n}} \leq \frac{1}{(1 + \eta)^{2n}}.$$

Now since $d_{\text{TV}}(X, Y) \leq \frac{1}{4n}$, we see that

$$\Pr \left[L_{f_Y} > (1 + \eta)^2 \sqrt{\frac{2e}{\pi}} \right] \leq \frac{1}{(1 + \eta)^{2n}} + \frac{1}{4n} \leq \frac{1}{2n} \quad (4.3.4)$$

for n large enough (η will be chosen to be constant). Hence after additionally conditioning on the complement of event 4.3.4, our success probability is at least $1 - \frac{1}{n}$.

At this point, letting $\mathbf{s} = Y$, we see that \mathbf{s} specifies a density f_s on K satisfying

$$L_{f_s} \leq (1 + \eta)^2 \sqrt{\frac{2e}{\pi}}.$$

Furthermore since $\mathbf{s} \in n(\text{conv}\{K - \mathbf{b}, \mathbf{b} - K\})^*$, $\mathbf{b}(K) - \mathbf{b} \in \frac{1}{n+1}E_K$ and $E_K \subseteq K - \mathbf{b}$, we have that

$$\frac{\sup_{\mathbf{x} \in K - \mathbf{b}} f_{\mathbf{s}}(\mathbf{x})}{f_{\mathbf{s}}(\mathbf{b}(K - \mathbf{b}))} = \sup_{\mathbf{x} \in K - \mathbf{b}} e^{\langle \mathbf{s}, \mathbf{x} - \mathbf{b}(K) - \mathbf{b} \rangle} = \sup_{\mathbf{x} \in K - \mathbf{b}} e^{\langle \mathbf{s}, \mathbf{x} \rangle + \langle -\mathbf{s}, \mathbf{b}(K) - \mathbf{b} \rangle} \leq e^{n+1}.$$

Hence by Lemma 4.3.8, letting $\sqrt{n}E(A) = T$, and $\delta = e^{\frac{1}{n}}$, we get that

$$N(K, \sqrt{n}E(A)) \leq (12\delta)^n \frac{4}{3} \frac{\sup_{\mathbf{x} \in K - \mathbf{b}} f_{\mathbf{s}}(\mathbf{x})}{f_{\mathbf{s}}(\mathbf{b}(K - \mathbf{b}))} \leq 12^n e \frac{4}{3} e^{n+1} \leq (12e(1 + \eta))^n$$

and

$$\begin{aligned} N(\sqrt{n}E(A), K) &\leq (12\delta^2)^n \text{vol}_n(\sqrt{n}B_2^n) \frac{4}{3} L_{f_{\mathbf{s}}}^n \\ &\leq 12^n e^2 (\sqrt{2\pi e}(1 + o(1)))^n \frac{4}{3} \left((1 + \eta)^3 \sqrt{2} \right)^n \leq (24e(1 + \eta)^3)^n \end{aligned}$$

for n large enough. Choosing $\eta > 0$ such that $(1 + \eta)^3 = \frac{25}{24}$ yields the result. \square

4.3.1 A Las Vegas Algorithm for Generating an M-Ellipsoid

The main result of this section is a $2^{O(n)}$ time Las Vegas algorithm to generate an M-Ellipsoid of any convex body. In the previous section, we showed that Klartag's construction yields an algorithm which succeeds with high probability. Here we show how to certify the outputted ellipsoid, removing the uncertainty of the Monte Carlo guarantee of Algorithm 4.4 (M-Gen).

To check whether that the candidate ellipsoid produced by M-Gen is an M-Ellipsoid, we use the Ellipsoid-Cover Algorithm (section 4.2) to check that both $N(K, E), N((K - K)^*, E^*) = 2^{O(n)}$ by constructing explicit coverings (if possible). Because $N(E, K) \approx N((K - K)^*, E^*)$ (up to $2^{\Theta(n)}$ factors) by the duality of entropy (Theorem 4.3.7), bounds on the size of such coverings suffice to prove that E is an M-Ellipsoid for K .

Theorem 4.3.3. *Algorithm M-Ellipsoid outputs (for n large enough) an ellipsoid $E \subseteq \mathbb{R}^n$ satisfying*

$$N(K, E) \leq \left(40e\sqrt{\frac{\pi e}{2}} \right)^n \quad \text{and} \quad N(E, K) \leq \left(901e\sqrt{\frac{\pi e}{2}} \cdot 289 \right)^n \quad (4.3.5)$$

Algorithm 4.5 M-Ellipsoid-Vegas: Generates a guaranteed M-Ellipsoid

Input: A weak membership oracle O_K for a (\mathbf{a}_0, r, R) -centered convex body K .

Output: An M-Ellipsoid E of K .

- 1: Generate a candidate M-Ellipsoid $E = E(A)$ of K using M-Gen (Algorithm 4.4).
If M-Gen fails, restart.
 - 2: $size \leftarrow 0$.
 - 3: **for all** $\mathbf{c} \in \text{Ellipsoid-Cover}(K, E)$ **do**
 - 4: $size \leftarrow size + 1$.
 - 5: **if** $size > (40e\sqrt{\frac{\pi e}{2}})^n$ **then**
 - 6: restart.
 - 7: $size \leftarrow 0$.
 - 8: **for all** $\mathbf{c} \in \text{Ellipsoid-Cover}((K - K)^*, E^*)$ **do**
 - 9: $size \leftarrow size + 1$.
 - 10: **if** $size > (901e\sqrt{\frac{\pi e}{2}})^n$ **then**
 - 11: restart.
 - 12: **return** E .
-

in expected time $(901e\sqrt{\frac{\pi e}{2}})^n \cdot \text{poly}(\cdot)$ using polynomial space.

Proof. The algorithm proceeds by first generating a candidate M-Ellipsoid E for K using M-Gen. Following this, it verifies that the covering numbers $N(K, E)$ and $N(E, K)$ are not too large by using the algorithm Ellipsoid-Cover. If any of the verification steps fails, the algorithm is restarted from the beginning.

Verifying the covering numbers. To verify the covering number $N(K, E)$, we call Algorithm Ellipsoid-Cover(K, E) and simply count to make sure the covering is not too large. To verify the covering number $N(E, K)$, we apply the same procedure however on the body $(K - K)^*$ and ellipsoid $E^* = E(A^{-1})$. To call Ellipsoid-Cover on $(K - K)^*$ and $E^* = E(A^{-1})$, we need to construct a weak membership oracle for $(K - K)^*$. First, from the guarantees on K , i.e. $rB_2^n \subseteq K - \mathbf{a}_0 \subseteq RB_2^n$, we have that $\frac{1}{2R}B_2^n \subseteq (K - K)^* \subseteq \frac{1}{2r}B_2^n$, i.e. $(K - K)^*$ is $(0, r\frac{1}{2R}, \frac{1}{2r})$ -centered. To build a weak membership oracle for $(K - K)^*$ we use the following characterization:

$$\mathbf{v} \in (K - K)^* \Leftrightarrow \sup_{\mathbf{x} \in K} \langle \mathbf{v}, \mathbf{x} \rangle - \inf_{\mathbf{x} \in K} \langle \mathbf{v}, \mathbf{x} \rangle \leq 1.$$

From the above, we will be able to build a weak membership oracle for $(K - K)^*$ by approximately maximizing and minimizing with respect to \mathbf{v} over K . This can readily be done via the ellipsoid algorithm (see Theorem 2.5.9). The exact details of the oracle construction are almost identical to the construction of the membership oracle for the polar in Algorithm 4.4; we leave the full analysis as an exercise to the reader.

Correctness: We must show that if the algorithm succeeds, the returned ellipsoid E satisfies

$$N(K, E) \leq \left(40e\sqrt{\frac{\pi e}{2}}\right)^n \quad N(E, K) \leq \left(901e\sqrt{\frac{\pi e}{2}} \cdot 289\right)^n$$

These guarantees depend only on the correctness of the algorithm Ellipsoid-Cover. The first counting test success if and only if Ellipsoid-Cover(K, E) returns a cover of size at most $(40e\sqrt{\frac{\pi e}{2}})^n$. Upon successful termination, the first requirement is therefore met. For the second test, we check that the cover produced by Ellipsoid-Cover($(K - K)^*, E^*$) is smaller than $(901\sqrt{\frac{\pi e}{2}})^n$. Now by Theorem 4.3.7, since E^* is centrally symmetric, we have that

$$N(E, K) \leq 289^n N((K - K)^*, E^*) \leq \left(901\sqrt{\frac{\pi e}{2}} \cdot 289\right)^n$$

for n large enough. Therefore the second requirement is also met.

Runtime: The main contributions to the running time comes the covering verification steps (i.e. checking that $N(K, E)$ and $N(E, K)$ are not too large). Since we halt the enumeration of the coverings if any of them grows larger than $(901e\sqrt{\frac{\pi e}{2}})^n$, none of the inner loops execute more than this number of times. Therefore the runtime of a single iteration of the main loop takes at most $(901e\sqrt{\frac{\pi e}{2}})^n \text{poly}(\cdot)$. Furthermore, all the algorithms invoked during the main loop require at most polynomial space. Therefore, to complete the runtime analysis, it suffices to show that the main loop is executes $O(1)$ times on expectation.

To begin this analysis, we condition the run of the main loop on the event that M-Gen returns an ellipsoid E satisfying

$$N(K, E) \leq (13e)^n \quad \text{and} \quad N(E, K) \leq (25e)^n. \quad (4.3.6)$$

By the guarantees on M-Gen, the probability of satisfying (4.3.6) is at least $1 - \frac{1}{n}$. Now we examine the covering number verification step. By the guarantees on Algorithm Ellipsoid-Cover and the conditioning (4.3.6), the size of the covering generated of K by E has size at most

$$\left(3\sqrt{\frac{\pi e}{2}}(1 + o(1))\right)^n N(K, E) \leq \left(3\sqrt{\frac{\pi e}{2}} \cdot 13e(1 + o(1))\right)^n \leq \left(40e\sqrt{\frac{\pi e}{2}}\right)^n,$$

for n large enough. Hence (for n large enough), the counting test for $N(K, E)$ is guaranteed to pass. Since E is centrally symmetric, by Theorem 4.3.7, we have that $N((K - K)^*, E^*) \leq (12(1 + o(1)))^n N(E, K)$. Therefore by the guarantees on Ellipsoid-Cover, the cover generate of $(K - K)^*$ by E^* has size at most

$$\begin{aligned} \left(3\sqrt{\frac{\pi e}{2}}(1 + o(1))\right)^n N((K - K)^*, E^*) &\leq \left(3\sqrt{\frac{\pi e}{2}} \cdot 12(1 + o(1))\right)^n N(K, E) \\ &\leq \left(3\sqrt{\frac{\pi e}{2}} \cdot 12 \cdot 25e(1 + o(1))\right)^n \leq \left(901e\sqrt{\frac{\pi e}{2}}\right)^n \end{aligned}$$

for n large enough. Therefore, the counting test for $N((K - K)^*, E^*)$ is also guaranteed to pass.

Finally, we get that the probability that each execution of the loop terminates successfully is at least $1 - \frac{1}{n}$. The expected number of runs of the loop is therefore $O(1)$ as needed. \square

4.3.2 Helper Algorithms

In this section, we provide some technical helper algorithms for the algorithms in the previous sections. These are straightforward applications of the fundamental algorithms presented in the section 2.5.3.

Corollary 4.3.4 (Algorithm Estimate-Covariance). *Let $K \subseteq \mathbb{R}^n$ be an (a_0, r, R) -centered convex body given by a weak membership oracle O_K . Let $f : K \rightarrow \mathbb{R}_+$ be a polynomial time computable log-concave function satisfying*

$$\sup_{\mathbf{x} \in K} f(\mathbf{x}) \leq e^{2n} f(0).$$

Then an ellipsoid $E(A)$, $A \in \mathbb{Q}^{n \times n}$, can be computed satisfying

$$e^{-\frac{1}{n}} E_f \subseteq E(A) \subseteq e^{\frac{1}{n}} E_f$$

with probability $1 - \delta$ in polynomial time.

Proof. Using Theorem 2.5.11, we can compute a matrix $B \subseteq \mathbb{Q}^{n \times n}$ satisfying

$$|\mathbf{x}^t (B - \text{cov}(f)) \mathbf{x}| \leq \frac{1}{n} \mathbf{x}^t \text{cov}(f) \mathbf{x} \quad \forall \mathbf{x} \in \mathbb{R}^n, \quad (4.3.7)$$

with probability $1 - \delta$ in polynomial time. We now condition on the event (4.3.7). Remembering that $\mathbf{x}^t B \mathbf{x} = \|\mathbf{x}\|_B^2$ and $\mathbf{x}^t \text{cov}(f) \mathbf{x} = \|\mathbf{x}\|_{\text{cov}(f)}^2$, we may rewrite (4.3.7) as

$$\sqrt{\frac{n-1}{n}} \|\mathbf{x}\|_{\text{cov}(f)} \leq \|\mathbf{x}\|_B \leq \sqrt{\frac{n+1}{n}} \|\mathbf{x}\|_{\text{cov}(f)}$$

From the above, we see that the ellipsoid $E(\text{cov}(f)) = \{\mathbf{x} : \|\mathbf{x}\|_{\text{cov}(f)} \leq 1\}$ and $E(B) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_B \leq 1\}$ satisfy

$$\sqrt{\frac{n}{n+1}} E(\text{cov}(f)) \subseteq E(B) \subseteq \sqrt{\frac{n}{n-1}} E(\text{cov}(f)) \quad (4.3.8)$$

Remembering that the polar ellipsoids satisfy

$$E(B)^* = E(B^{-1}) \quad \text{and} \quad E(\text{cov}(f))^* = E(\text{cov}(f)^{-1}) = E_f.$$

where the last equality follows by the definition of E_f . Taking the polars of the above ellipsoids, the containment relationships in (4.3.8) flip, and we get

$$\sqrt{\frac{n-1}{n}} E_f \subseteq E(B^{-1}) \subseteq \sqrt{\frac{n+1}{n}} E_f \quad (4.3.9)$$

Now using the inequalities $1 - \frac{1}{n} \geq e^{-\frac{2}{n}}$ for $n \geq 3$ and $1 + \frac{1}{n} \leq e^{\frac{2}{n}}$, we see that (4.3.9) implies

$$e^{-\frac{1}{n}} E_f \subseteq E(B^{-1}) \subseteq e^{\frac{1}{n}} E_f$$

as needed. Letting $A = B^{-1}$, the ellipsoid $E(A)$ satisfies the desired requirements. \square

Corollary 4.3.5 (Algorithm Estimate-Centroid). *There is a probabilistic algorithm Estimate-Centroid that, given a (\mathbf{a}_0, r, R) -centered convex body K presented by a weak membership oracle O_K and some $\delta > 0$, in polynomial time either outputs FAIL (with probability at most δ) or some $\mathbf{b} \in K$ such that:*

$$\mathbf{b} + \frac{r}{2(n+1)\sqrt{n}} B_2^n \subseteq K \subseteq \mathbf{b} + 2RB_2^n$$

where with probability at least $1 - \delta$,

$$\mathbf{b} - \mathbf{b}(K) \in \frac{1}{n+1} E_K.$$

Proof. Using Theorem 2.5.11, we compute a center $\mathbf{b} \in K$ satisfying

$$|\langle \mathbf{x}, \mathbf{b} - \mathbf{b}(K) \rangle| \leq \frac{1}{(n+1)^2} \mathbf{x}^t \text{cov}(K) \mathbf{x} \quad \forall \mathbf{x} \in \mathbb{R}^n, \quad (4.3.10)$$

with probability $1 - \delta$ in polynomial time

First, check whether

$$O_K \left(\mathbf{b} \pm \frac{3r}{4(n+1)} \mathbf{e}_i, \frac{r}{4(n+1)\sqrt{n}} \right) = 1 \quad \text{for } 1 \leq i \leq n \quad (4.3.11)$$

If any of the above tests fail, abort and return FAIL.

Let $\delta = \frac{r}{n+1}$. If these tests pass, by the properties of O_K we know that

$$\mathbf{b} + \frac{3\delta}{4} \text{conv}\{\pm \mathbf{e}_1, \dots, \pm \mathbf{e}_n\} \subseteq K^{\frac{\delta}{4\sqrt{n}}} \Rightarrow \mathbf{b} + \frac{3\delta}{4\sqrt{n}} B_2^n \subseteq K^{\frac{\delta}{4\sqrt{n}}} \Rightarrow \mathbf{b} + \frac{\delta}{2\sqrt{n}} B_2^n \subseteq K$$

Since $\mathbf{b} \in K \subseteq RB_2^n$, we clearly also have that $K \subseteq \mathbf{b} + 2RB_2^n$. Hence conditioned on outputting \mathbf{b} , we have that

$$\mathbf{b} + \frac{r}{2(n+1)\sqrt{n}} B_2^n \subseteq K \subseteq \mathbf{b} + 2RB_2^n$$

as needed.

We now show that if the event (4.3.10) holds, then the above test will pass and condition (b) will also be satisfied. Since this event holds with probability $1 - \delta$, this will suffice to prove the statement.

For the center \mathbf{b} , we note that for all $\mathbf{x} \in (n+1)E(\text{cov}(K))$, by equation (4.3.10) we have that

$$|\langle \mathbf{b} - \mathbf{b}(K), \mathbf{x} \rangle| \leq \frac{1}{(n+1)^2} \mathbf{x}^t \text{cov}(K) \mathbf{x} \leq \frac{1}{(n+1)^2} (n+1)^2 = 1$$

Therefore, we have that $\mathbf{b} - \mathbf{b}(K) \in ((n+1)E(\text{cov}(K)))^* = \frac{1}{n+1}E_K$ as needed.

We now show that the tests must all pass. From Theorem 2.3.5, we know that

$$\mathbf{b}(K) + \sqrt{\frac{n+2}{n}}E_K \subseteq K \subseteq \mathbf{b}(K) + \sqrt{n(n+2)}E_K$$

By the guarantee on O_K , we know that $rB_2^n \subseteq \mathbf{b}(K) + \sqrt{n(n+2)}E_K$. But we have that

$$\begin{aligned} rB_2^n - \mathbf{b}(K) &\subseteq \sqrt{n(n+2)}E_K \Rightarrow rB_2^n + \mathbf{b}(K) \subseteq \sqrt{n(n+2)}E_K \\ &\Rightarrow \frac{1}{2}(rB_2^n - \mathbf{b}(K)) + \frac{1}{2}(rB_2^n + \mathbf{b}(K)) \subseteq \sqrt{n(n+2)}E_K \\ &\Rightarrow rB_2^n \subseteq \sqrt{n(n+2)}E_K \end{aligned}$$

since both E_K and B_2^n are symmetric. From the inequality $n+1 \geq \sqrt{n(n+2)}$, we have that

$$\frac{r}{n+1}B_2^n \subseteq \frac{\sqrt{n(n+2)}}{n+1}E_K \subseteq E_K \quad (4.3.12)$$

Since $\mathbf{b} - \mathbf{b}(K) \in \frac{1}{n+1}E_K$ by assumption, and $\sqrt{\frac{n+2}{n}}E_K + \mathbf{b}(K) \subseteq K$, we get that

$$\mathbf{b} \in \mathbf{b}(K) + \frac{1}{n+1}E_K \Rightarrow \mathbf{b} + E_K \subseteq \mathbf{b}(K) + \frac{n+2}{n+1}E_K \Rightarrow \mathbf{b} + E_K \subseteq \mathbf{b}(K) + \sqrt{\frac{n+2}{n}}E_K \subseteq K$$

Therefore by 4.3.12 we have that $\mathbf{b} + \frac{r}{n+1}B_2^n \subseteq K$. Letting $\delta = \frac{r}{n+1}$, from the previous sentence we see that

$$\mathbf{b} \pm \frac{3}{4}\delta \mathbf{e}_i \in K^{-\frac{\delta}{4}} \subseteq K^{-\frac{\delta}{4\sqrt{n}}}$$

Therefore by the properties of O_K , the tests in 4.3.11 must all pass. The claim thus holds. \square

4.3.3 Geometric Estimates

Here we list and prove the necessary geometric inequalities that are needed in the previous sections. We begin with a slight extension of Theorem 2.3.7.

Theorem 4.3.6. *Let K be a convex body such that $\mathbf{b}(K) \in tE_K$, for some $t \in [0, 1)$.*

Then

$$\text{vol}_n(K \cap -K) \geq \left(\frac{1-t}{2}\right)^n \text{vol}_n(K)$$

Proof. From Theorem 2.3.7 we have that

$$\frac{1}{2^n} \text{vol}_n(K) \leq \text{vol}_n(K - \mathbf{b}(K) \cap -K + \mathbf{b}(K)) = \text{vol}_n(K \cap -K + 2\mathbf{b}(K))$$

Next, we note that for $\mathbf{x} \in \mathbb{R}^n$

$$K \cap (-K + 2\mathbf{x}) \neq \emptyset \Leftrightarrow 2\mathbf{x} \in K + K \Leftrightarrow \mathbf{x} \in K \quad (4.3.13)$$

Since $\mathbf{b}(K) \in tE_K$ and $\mathbf{b}(K) + E_K \subseteq K$, we see that $(1-t)E_K \subseteq K$. Hence we can write

$$0 = t(-2n\mathbf{b}(K)) + (1-t)2\mathbf{b}(K),$$

where $-n\mathbf{b}(K) \in -(1-t)E_K = (1-t)E_K \subseteq K$. Now we see that

$$t(K \cap (-K + -2n\mathbf{b}(K))) + (1-t)(K \cap (-K + 2\mathbf{b}(K))) \subseteq K \cap -K$$

where both sets on the left hand side are non-empty by (4.3.13). Therefore by the Brunn-Minkowski inequality, we have that

$$\begin{aligned} \text{vol}_n(K \cap -K)^{\frac{1}{n}} &\geq t \text{vol}_n(K \cap (-K + -2n\mathbf{b}(K)))^{\frac{1}{n}} + (1-t) \text{vol}_n(K \cap (-K + 2\mathbf{b}(K)))^{\frac{1}{n}} \\ &\geq (1-t) \text{vol}_n(K \cap (-K + 2\mathbf{b}(K)))^{\frac{1}{n}} \geq \frac{1-t}{2} \text{vol}_n(K)^{\frac{1}{n}} \end{aligned}$$

Therefore we get that

$$\text{vol}_n(K \cap -K) \geq \left(\frac{1-t}{2}\right)^n \text{vol}_n(K)$$

as needed. □

The next lemma is a slight specialization of [95, Theorem 5]. We require this inequality for the M-Ellipsoid certification procedure.

Theorem 4.3.7 (Duality of Entropy). *Let $K, T \subseteq \mathbb{R}^n$ be convex bodies where T is centrally symmetric. Then*

$$N(T, K) \leq ((1 + o(1))288)^n \cdot N((K - K)^*, T^*)$$

and

$$N((K - K)^*, T^*) \leq (12(1 + o(1)))^n \cdot N(T, K).$$

Proof. Since the above quantities are invariant under shifts of K , we may shift K so that $\mathbf{b}(K) = 0$. Applying Theorem 2.3.7, we see that $\text{vol}_n(K - K) \leq 4^n \text{vol}_n(K) \leq 8^n \text{vol}_n(K \cap -K)$, where we note that since $\mathbf{0} \in K$ we have that $K \cap -K \subseteq K \subseteq K - K$. Next applying the covering estimates from Lemma 2.3.9, we get that

$$N(K - K, K) \leq N(K - K, K \cap -K) \leq 3^n \frac{\text{vol}_n(K - K)}{\text{vol}_n(K \cap -K)} \leq 24^n.$$

From here, we see that

$$N(T, K) \leq N(T, K - K)N(K - K, K) \leq 24^n N(T, K - K).$$

Next since both T and $K - K$ are centrally symmetric, we apply Lemma 2.3.9 to get that

$$N(T, (K - K)) \leq 3^n \frac{\text{vol}_n(T)}{\text{vol}_n((K - K) \cap T)}.$$

Now we note that $((K - K) \cap T)^* = \text{conv}\{(K - K)^*, T^*\}$. Hence applying the Blaschke-Santaló inequality to $\text{vol}_n(T)$ and the Bourgain-Milman inequality to $\text{vol}_n((K - K) \cap T)$ we get that

$$3^n \frac{\text{vol}_n(T)}{\text{vol}_n((K - K) \cap T)} \leq (6(1 + o(1)))^n \frac{\text{vol}_n(\text{conv}\{(K - K)^*, T^*\})}{\text{vol}_n(T^*)}$$

Since 0 is both in $(K - K)^*$ and T^* , we see that $\text{conv}\{(K - K)^*, T^*\} \subseteq (K - K)^* + T^*$ and hence

$$(6(1 + o(1)))^n \frac{\text{vol}_n(\text{conv}\{(K - K)^*, T^*\})}{\text{vol}_n(T^*)} \leq (6(1 + o(1)))^n \frac{\text{vol}_n((K - K)^* + T^*)}{\text{vol}_n(T^*)}.$$

Lastly, applying Lemma 2.3.9 to the last estimate, we get that

$$(6(1 + o(1)))^n \frac{\text{vol}_n((K - K)^* + T^*)}{\text{vol}_n(T^*)} \leq (12(1 + o(1)))^n N((K - K)^*, T^*).$$

Combining the above estimates yields the first desired inequality.

Now switching the roles $(K - K)$ and T with $(K - K)^*$ and T^* , we have that

$$N((K - K)^*, T^*) \leq (12(1 + o(1)))^n N(T, K - K) \leq (12(1 + o(1)))^n N(T, K),$$

yielding the second inequality. \square

We now make precise the relationship between the isotropic constant of the exponential reweightings defined by Klartag [78] and the M-Ellipsoid.

Lemma 4.3.8. *Let $K \subseteq \mathbb{R}^n$ be a convex body. Take $\mathbf{s} \in \mathbb{R}^n$ and let $f_{\mathbf{s}}(\mathbf{x}) = e^{(\mathbf{s}, \mathbf{x})}$ for $\mathbf{x} \in K$ and 0 otherwise. Let $T \subseteq \mathbb{R}^n$ be a convex body such that for some $\delta \geq 1$ we have that*

$$\frac{\sqrt{n}}{\delta} E_{f_{\mathbf{s}}} \subseteq T \subseteq \delta \sqrt{n} E_{f_{\mathbf{s}}} \quad (4.3.14)$$

where $E_{f_{\mathbf{s}}}$ is the inertial ellipsoid of $f_{\mathbf{s}}$. Then we have that

$$N(K, T) \leq (12\delta)^n \frac{4}{3} \frac{\sup_{\mathbf{x} \in K} f_{\mathbf{s}}(\mathbf{x})}{f_{\mathbf{s}}(\mathbf{b}(K))} \quad \text{and} \quad N(T, K) \leq (12\delta^2)^n \text{vol}_n(\sqrt{n} B_2^n) \frac{4}{3} L_{f_{\mathbf{s}}}^n \quad (4.3.15)$$

where $\mathbf{b}(K)$ is the centroid of K , and $L_{f_{\mathbf{s}}}$ is the isotropic constant of $f_{\mathbf{s}}$.

Proof. Since the above estimates are all invariant under shifts of K , we may assume that $\mathbf{b}(f_{\mathbf{s}}) = \mathbf{0}$ (centroid of $f_{\mathbf{s}}$). We note that $\mathbf{b}(f_{\mathbf{s}}) \in K$ always and hence $\mathbf{0} \in K$. Let X be distributed as $\pi_{f_{\mathbf{s}}}$, where $\pi_{f_{\mathbf{s}}}$ is the probability measure induced by $f_{\mathbf{s}}$. So we have that $\mathbb{E}[X] = \mathbf{b}(f_{\mathbf{s}}) = \mathbf{0}$ and $\mathbb{E}[XX^t] = \text{cov}(f_{\mathbf{s}})$.

Remember that $E_{f_{\mathbf{s}}} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^t \text{cov}(f_{\mathbf{s}})^{-1} \mathbf{x} \leq 1\}$,

therefore $\|\mathbf{x}\|_{E_{f_{\mathbf{s}}}} = \sqrt{\mathbf{x}^t \text{cov}(f_{\mathbf{s}})^{-1} \mathbf{x}}$. Now note that

$$\begin{aligned} \mathbb{E}[\|X\|_{E_{f_{\mathbf{s}}}}^2] &= \mathbb{E}[X^t \text{cov}(f_{\mathbf{s}})^{-1} X] = \mathbb{E}[\text{trace}[\text{cov}(f_{\mathbf{s}})^{-1} X X^t]] = \text{trace}[\text{cov}(f_{\mathbf{s}})^{-1} \mathbb{E}[X X^t]] \\ &= \text{trace}[\text{cov}(f_{\mathbf{s}})^{-1} \text{cov}(f_{\mathbf{s}})] = \text{trace}[\text{Id}_n] = n. \end{aligned}$$

Now by Markov's inequality, we have that

$$\pi_{f_s}(2\sqrt{n}E_{f_s}) = 1 - \Pr[\|X\|_{E_{f_s}} > 2\sqrt{n}] \geq 1 - \frac{\mathbb{E}[\|X\|_{E_{f_s}}^2]}{4n} = 1 - \frac{n}{4n} = \frac{3}{4}. \quad (4.3.16)$$

By Jensen's inequality, we see that

$$\begin{aligned} \int_K f_s(\mathbf{x})dx &= \int_K e^{\langle s, \mathbf{x} \rangle} dx = \text{vol}_n(K) \int_K \frac{e^{\langle s, \mathbf{x} \rangle}}{\text{vol}_n(K)} dx \geq \text{vol}_n(K) e^{\langle s, \mathbf{b}(K) \rangle} \\ &= \text{vol}_n(K) f_s(\mathbf{b}(K)), \end{aligned} \quad (4.3.17)$$

where $\mathbf{b}(K)$ is the centroid of K .

Using (4.3.17) and (4.3.16) we see that

$$\text{vol}_n(2\sqrt{n}E_{f_s} \cap K) \geq \frac{\int_{2\sqrt{n}E_{f_s}} f_s(\mathbf{x})dx}{\sup_{\mathbf{x} \in K} f_s(\mathbf{x})} \geq \frac{3}{4} \frac{\int_K f_s(\mathbf{x})dx}{\sup_{\mathbf{x} \in K} f_s(\mathbf{x})} \geq \frac{3}{4} \frac{f_s(\mathbf{b}(K))}{\sup_{\mathbf{x} \in K} f_s(\mathbf{x})} \text{vol}_n(K). \quad (4.3.18)$$

Using that $\frac{\sqrt{n}}{\delta}E_{f_s} \subseteq T$, $0 \in K$, $\delta \geq 1$, and by (4.3.18) we get that

$$\begin{aligned} \text{vol}_n(T \cap K) &\geq \text{vol}_n\left(\frac{\sqrt{n}}{\delta}E_{f_s} \cap K\right) = \left(\frac{1}{\delta}\right)^n \text{vol}_n(\sqrt{n}E_{f_s} \cap \delta K) \\ &\geq \left(\frac{1}{\delta}\right)^n \text{vol}_n\left(\sqrt{n}E_{f_s} \cap \frac{1}{2}K\right) = \left(\frac{1}{2\delta}\right)^n \text{vol}_n(2\sqrt{n}E_{f_s} \cap K) \\ &\geq \left(\frac{1}{2\delta}\right)^n \frac{3}{4} \frac{f_s(\mathbf{b}(K))}{\sup_{\mathbf{x} \in K} f_s(\mathbf{x})} \text{vol}_n(K). \end{aligned} \quad (4.3.19)$$

Using the definition of L_{f_s} , (4.3.16), $\sqrt{n}E_{f_s} \subseteq \delta T$ and that $0 \in K$, we get that

$$\begin{aligned} \det(\text{cov}(f_s))^{\frac{1}{2}} &= L_K^n \frac{\int_K f_s(\mathbf{x})dx}{\sup_{\mathbf{x} \in K} f_s(\mathbf{x})} \leq L_K^n \frac{4}{3} \frac{\int_{2\sqrt{n}E_{f_s}} f_s(\mathbf{x})dx}{\sup_{\mathbf{x} \in K} f_s(\mathbf{x})} \\ &\leq L_K^n \frac{4}{3} \text{vol}_n(2\sqrt{n}E_{f_s} \cap K) \leq L_K^n \frac{4}{3} \text{vol}_n(2\delta T \cap K) \\ &\leq (2\delta L_K)^n \frac{4}{3} \text{vol}_n(T \cap K). \end{aligned} \quad (4.3.20)$$

Using that $T \subseteq \delta\sqrt{n}E_{f_s}$ and the ellipsoid volume formula, we have that

$$\text{vol}_n(T) \leq \text{vol}_n(\delta\sqrt{n}E_{f_s}) = \delta^n \text{vol}_n(\sqrt{n}B_2^n) \det(\text{cov}(f_s))^{\frac{1}{2}}. \quad (4.3.21)$$

Combining equations (4.3.20),(4.3.21) we get that

$$\text{vol}_n(T) \leq (2\delta^2 L_K)^n \text{vol}_n(\sqrt{n}B_2^n) \frac{4}{3} \text{vol}_n(T \cap K). \quad (4.3.22)$$

Now applying Lemma 2.3.9 to the inequalities (4.3.19),(4.3.22) the theorem follows. \square

From Lemma 4.3.8, we see that if the slicing conjecture is true, then for any convex body, its inertial ellipsoid appropriately scaled is an M -ellipsoid. To bypass this, Klartag shows that for any convex body K , there exists a “mild” exponential reweighting $f_{\mathbf{s}}$ of the uniform density on K with bounded isotropic constant. As one can see from Lemma 4.3.8, the severity of the reweighting controls $N(K, \sqrt{n}E_{f_{\mathbf{s}}})$ whereas the isotropic constant of $f_{\mathbf{s}}$ controls $N(\sqrt{n}E_{f_{\mathbf{s}}}, K)$.

The main tool to establish the existence of “good” exponential reweightings for K is the following lemma, which one can extract from the proof of Theorem 4.3.1 in [78]. We will use it here for $\epsilon = 1$, in which case the expectation below is of order $2^{O(n)}$. The argument is essentially identical to that of [78]; we include it for completeness.

Theorem 4.3.9 ([78]). *Let $K \subseteq \mathbb{R}^n$ be a convex body such that $\mathbf{b}(K) \in \frac{1}{n+1}E_K$. For $\mathbf{s} \in \mathbb{R}^n$, let $f_{\mathbf{s}} : K \rightarrow \mathbb{R}^+$ denote the function $f_{\mathbf{s}}(\mathbf{x}) = e^{\langle \mathbf{s}, \mathbf{x} \rangle}$, $\mathbf{x} \in K$. Let X be distributed as $\epsilon n (\text{conv}\{K, -K\})^*$ for some real $\epsilon > 0$. Then we have*

$$\mathbb{E}[L_{f_X}^{2n}] \leq \left((1 + o(1)) \sqrt{\frac{2}{\pi e}} \frac{e^\epsilon}{\sqrt{\epsilon}} \right)^{2n}$$

Proof. For $\mathbf{s} \in \mathbb{R}^n$ define $f_{\mathbf{s}} : K \rightarrow \mathbb{R}_+$ by $f_{\mathbf{s}}(\mathbf{x}) = e^{\langle \mathbf{s}, \mathbf{x} \rangle}$ for $\mathbf{x} \in K$. In Lemma 3.2 of [78] is it shown that

$$\int_{\mathbb{R}^n} \det(\text{cov}(f_{\mathbf{s}})) d\mathbf{s} = \text{vol}_n(K) \quad (4.3.23)$$

By Theorem 2.3.5, we have that $E_K + \mathbf{b}(K) \subseteq K$. Since $\mathbf{b}(K) \in \frac{1}{n+1}E_K$ by assumption, we see that $\frac{n}{n+1}E_K \subseteq E_K + \mathbf{b}(K) \subseteq K$. Hence $\mathbf{0} \in K$. From [112], we know that for any convex body K such that $\mathbf{0} \in K$, we have that $\text{vol}_n(\text{conv}\{K, -K\}) \leq 2^n \text{vol}_n(K)$.

Let $L = \text{conv}\{K, -K\}$. Note that

$$L^* = (\text{conv}\{K, -K\})^* = \{\mathbf{y} \in \mathbb{R}^n : |\langle \mathbf{x}, \mathbf{y} \rangle| \leq 1, \forall \mathbf{x} \in K\}$$

Since L is centrally symmetric by the Bourgain-Milman inequality (Theorem 2.3.6), we have that

$$\text{vol}_n(L^*) \text{vol}_n(L) \geq \left((1 + o(1)) \frac{\pi e}{n} \right)^n$$

Hence we get that

$$\text{vol}_n(L^*) \geq \left(\frac{(1+o(1))\pi e}{n \text{vol}_n(L)^{\frac{1}{n}}} \right)^n \geq \left(\frac{(1+o(1))\pi e}{2n \text{vol}_n(K)^{\frac{1}{n}}} \right)^n \quad (4.3.24)$$

Take $\mathbf{s} \in \epsilon n L^*$. We examine the properties of $f_{\mathbf{s}} : K \rightarrow \mathbb{R}_+$. Since $\mathbf{s} \in \epsilon n L^*$, we see that

$$\sup_{\mathbf{x} \in K} f_{\mathbf{s}}(\mathbf{x}) = e^{\sup_{\mathbf{x} \in K} \langle \mathbf{s}, \mathbf{x} \rangle} \leq e^{\epsilon n} \quad (4.3.25)$$

Since $\mathbf{b}(K) \subseteq \frac{1}{n+1} E_K \subseteq \frac{1}{n} K$ and $\mathbf{s} \in \epsilon n (\text{conv}\{K, -K\})^*$, we see that $|\langle \mathbf{s}, \mathbf{b}(K) \rangle| \leq \epsilon$.

Now by Jensen's inequality, we have that

$$\begin{aligned} \int_K e^{\langle \mathbf{s}, \mathbf{x} \rangle} dx &= \text{vol}_n(K) \left(\int_K e^{\langle \mathbf{s}, \mathbf{x} \rangle} \frac{dx}{\text{vol}_n(K)} \right) \geq \text{vol}_n(K) e^{\int_K \langle \mathbf{s}, \mathbf{x} \rangle \frac{dx}{\text{vol}_n(K)}} \\ &= \text{vol}_n(K) e^{\langle \mathbf{s}, \mathbf{b}(K) \rangle} \geq \text{vol}_n(K) e^{-\epsilon} \end{aligned}$$

Now we see that

$$\begin{aligned} L_{f_{\mathbf{s}}}^{2n} &= \left(\sup_{\mathbf{x} \in K} \frac{f_{\mathbf{s}}(\mathbf{x})}{\int_K f_{\mathbf{s}}(\mathbf{x}) dx} \right)^2 \det(\text{cov}(f_{\mathbf{s}})) \leq \left(\frac{e^{\epsilon n}}{\text{vol}_n(K) e^{-\epsilon}} \right)^2 \det(\text{cov}(f_{\mathbf{s}})) \\ &= \frac{e^{2(n+1)\epsilon}}{\text{vol}_n(K)^2} \det(\text{cov}(f_{\mathbf{s}})) \end{aligned} \quad (4.3.26)$$

Applying inequality (4.3.26), Lemma 3.2 of [78], and equation (4.3.24), we get that

$$\begin{aligned} \frac{1}{\text{vol}_n(\epsilon n L^*)} \int_{\epsilon n L^*} L_{f_{\mathbf{s}}}^{2n} ds &\leq \frac{e^{2(n+1)\epsilon}}{\text{vol}_n(\epsilon n L^*) \text{vol}_n(K)^2} \int_{\epsilon n L^*} \text{vol}_n(K)^2 \det(\text{cov}(f_{\mathbf{s}})) ds \\ &\leq \frac{e^{2(n+1)\epsilon}}{\text{vol}_n(\epsilon n L^*) \text{vol}_n(K)^2} \text{vol}_n(K) \leq \left(\frac{(1+o(1))e^{2\epsilon}}{\epsilon n \text{vol}_n(L^*)^{\frac{1}{n}} \text{vol}_n(K)^{\frac{1}{n}}} \right)^n \\ &\leq \left(\frac{(1+o(1))2e^{2\epsilon}}{\pi e \epsilon} \right)^n = \left((1+o(1)) \sqrt{\frac{2}{\pi e}} \frac{e^\epsilon}{\sqrt{\epsilon}} \right)^{2n} \end{aligned}$$

The above quantity is exactly $E[L_{f_X}]$ since X is uniform over $\epsilon n L^*$. The statement thus follows. \square

4.4 Milman's Construction

In this section, we provide a deterministic algorithmic implementation of Milman's M-Ellipsoid construction.

We begin with some background to his construction, and then proceed to outline its steps. Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body. To construct an M-Ellipsoid, a natural starting point is to try the the John (or Löwner) ellipsoid E of K , i.e. the maximum volume inscribed ellipsoid in K . Here in the worst case that $E \subseteq K \subseteq \sqrt{n}E$, and hence $\text{vol}_n(E) \leq \text{vol}_n(K) \leq n^{\frac{n}{2}}\text{vol}_n(E)$. From the classical bounds on the covering estimates (see Lemma 2.3.9), we remember that

$$\max\{N(K, E), N(E, K)\} = 2^{\Theta(n)} \frac{\max\{\text{vol}_n(E), \text{vol}_n(K)\}}{\text{vol}_n(E \cap K)}$$

Therefore, the covering estimates for the John ellipsoid are of order $n^{\Omega(n)}$ in the worst case. The main problem with the John ellipsoid is that the containment condition (i.e. $E \subseteq K$) is far too restrictive. More precisely, we may potentially be able to ensure a much larger intersection volume if we allow E to “stick out” of K a bit.

Moving in line with this intuition, the next ellipsoid which was considered is the ℓ -Ellipsoid of K (which we describe in section 4.4.1), which roughly corresponds to the largest “half-contained” ellipsoid in K . By half-contained we mean that $\text{vol}_n(E \cap K) \geq \frac{1}{2}\text{vol}_n(E)$. This condition is not imposed directly and is instead expressed (up to constant factor scaling of E) by imposing that a certain gaussian expectation related to K and E be small. Here a deep theorem of Pisier [105], shows that relaxing the containment condition in this way yields a drastic improvement in the size of the intersection. In particular, he shows that the ℓ -Ellipsoid E satisfies

$$\frac{\max\{\text{vol}_n(E), \text{vol}_n(K)\}}{\text{vol}_n(E \cap K)} = O(\log n)^n.$$

Furthermore, he shows that something stronger is true. Letting $D_K = d_{BM}(K, B_2^n)$, the Banach Mazur distance between K and the euclidean ball B_2^n (see Definition 2.3.1), the above ratio is bounded by $O(\log(D_K))^n$ (where the above estimate now follows from John’s theorem).

In the final development, Milman [92] discovered that ℓ -Ellipsoid construction can in fact be “amplified”, i.e. it can be used in an iterated fashion, via a process he calls

isomorphic symmetrization, to yield a sequence of ellipsoids with larger and larger intersection volumes with K . To understand this, we note that the main barrier to the ℓ -Ellipsoid having large intersection volume is D_K . What Milman noticed is that one apply a small surgery to K to yield a body K' , that has substantially smaller $D_{K'}$ and for which the ratio

$$\frac{\max\{\text{vol}_n(K), \text{vol}_n(K')\}}{\text{vol}_n(K \cap K')}$$

is not too large. The surgery here is simple, given an ℓ -Ellipsoid E of K , we let $K' = \text{conv}\{\frac{1}{C \log(D_K)}E, C \log(D_K)E \cap K\}$, for an appropriately chosen constant $C \geq 1$. From the surgery formula, it is easy to check the new Banach-Mazur distance satisfies $D_{K'} = O(\log(D_K)^2)$. Hence, after applying the above iteration i times, the Banach-Mazur distance of the resultant body to B_2^n drops like $\log^{(i)}(n)$, i.e. the iterated logarithm. Therefore after only $\log^*(n)$ iterations, the resultant body is in fact $O(1)$ -isomorphic to a ball and still contains large intersection volume with K . The ℓ -Ellipsoid for the final body now yields the desired M-Ellipsoid for K .

We devote the rest of this section to making the above outline precise and showing that each of the associated steps can be algorithmically implemented in a time and space efficient manner.

4.4.1 The Lewis Ellipsoid

Let α be a norm on $n \times n$ matrices. We define the dual norm α^* for any $S \in \mathbb{R}^{n \times n}$ as

$$\alpha^*(S) = \sup\{\text{tr}(SA) : A \in \mathbb{R}^{n \times n}, \alpha(A) \leq 1\}. \quad (4.4.1)$$

For a matrix $A \in \mathbb{R}^{n \times n}$, we denote its transpose by A^T , and its inverse (when it exists) by A^{-1} .

Theorem 4.4.1. [85] *For any norm α on $\mathbb{R}^{n \times n}$, there is an invertible linear transformation $A \in \mathbb{R}^{n \times n}$ such that*

$$\alpha(A) = 1 \text{ and } \alpha^*(A^{-1}) = n.$$

The proof of the above theorem is based on examining the properties of the optimal solution to the following mathematical program:

$$\begin{aligned}
& \max \det(A) \\
& \text{s.t.} \\
& A \in \mathbb{R}^{n \times n} \\
& \alpha(A) \leq 1
\end{aligned} \tag{4.4.2}$$

From here, showing that the optimal A satisfies $\alpha^*(A^{-1}) = n$ is a simple variational argument (reproduced in Lemma 4.4.11).

We will be interested in norms α of the following form. Let $K \subseteq \mathbb{R}^n$ denote a convex body, satisfying $\mathbf{0} \in \text{int}(K)$, with associated norm $\|\cdot\|_K$, and let γ_n denote the canonical Gaussian measure on \mathbb{R}^n (i.e. for $A \subseteq \mathbb{R}^n$ measurable $\gamma_n(A) = \frac{1}{\sqrt{2\pi}^n} \int_A e^{-\frac{1}{2}\|\mathbf{x}\|_2^2} d\mathbf{x}$). We define the ℓ -norm with respect to K for $A \in \mathbb{R}^{n \times n}$ as

$$\ell_K(A) = \left(\int \|A\mathbf{x}\|_K^2 d\gamma_n(\mathbf{x}) \right)^{1/2}$$

The ℓ -norm was first studied and defined by Tomczak-Jaegermann and Figiel [49].

The next crucial ingredient is a connection between the dual norm α^* defined above and the ℓ -norm with respect to the polar $K^* = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle \leq 1 \ \forall \mathbf{y} \in K\}$, namely,

$$\ell_{K^*}(A) = \left(\int \|A\mathbf{x}\|_{K^*}^2 d\gamma_n(\mathbf{x}) \right)^{1/2}.$$

Lemma 4.4.2. [106] *For $K \subseteq \mathbb{R}^n$ a symmetric convex body, and $A \in \mathbb{R}^{n \times n}$ we have that*

$$\ell_{K^*}(A^T) \leq 4(1 + \log_2 d_{BM}(K, B_2^n)) \ell_K^*(A)$$

where $d_{BM}(K, B_2^n)$ is the Banach-Mazur distance between K and B_2^n .

Combining Lemma 4.4.2 with the optimal solution to program 4.4.2 yields the following classical bound on the $\ell\ell^*$ estimate in convex geometry.

Theorem 4.4.3. *For $K \subseteq \mathbb{R}^n$ a symmetric convex body, we have that*

$$\begin{aligned} \ell\ell^*(K) &\stackrel{\text{def}}{=} \inf_{\substack{A \in \mathbb{R}^{n \times n} \\ \det(A)=1}} \ell_K(A)\ell_{K^*}(A^{-T}) \\ &\leq 4n(1 + \log_2 d_{BM}(K, B_2^n)) \leq 4n(1 + \frac{1}{2} \log_2 n) \end{aligned}$$

Note that the bound $\log_2 d_{BM}(K, B_2^n) \leq \frac{1}{2} \log_2 n$ is precisely John's Theorem. In the segway, the bound on $\ell\ell^*$ will yield the most important tool for the M-Ellipsoid construction.

4.4.2 Covering Numbers and Volume Estimates

Let $B_2^n \subseteq \mathbb{R}^n$ denote the n -dimensional Euclidean ball. Recall that $N(K, D)$ is the number of translates of D required to cover K . The following bounds for symmetric convex bodies $K, D \subset \mathbb{R}^n$ are classical.

Lemma 4.4.4. *For any two symmetric convex bodies K, D ,*

$$\frac{\text{vol}(K)}{\text{vol}(K \cap D)} \leq N(K, D) \leq 3^n \frac{\text{vol}(K)}{\text{vol}(K \cap D)}.$$

The next lemma is from [93].

Lemma 4.4.5. *Let $D \subseteq \alpha K$, $\alpha \geq 1$. Then,*

$$\text{vol}(\text{conv}\{K, D\}) \leq 4\alpha n N(D, K) \text{vol}(K).$$

The following are the Sudakov and dual Sudakov inequalities (see e.g., Section 6 of [53]).

Lemma 4.4.6 (Sudakov Inequality). *For any $t > 0$, and invertible matrix $A \in \mathbb{R}^{n \times n}$*

$$N(K, tAB_2^n) \leq e^{C\ell_{K^*}(A^{-T})^2/t^2}.$$

Lemma 4.4.7 (Dual Sudakov Inequality). *For any $t > 0$, and $A \in \mathbb{R}^{n \times n}$*

$$N(AB_2^n, tK) \leq e^{C\ell_K(A)^2/t^2}.$$

The following lemma gives a simple containment relationship (see e.g., [31]).

Lemma 4.4.8. *For any $A \in \mathbb{R}^{n \times n}$, A invertible, we have that*

$$\frac{1}{\ell_{K^*}(A^{-1})}K \subseteq AB_2^n \subseteq \ell_K(A)K$$

Proof. We first show that $E = AB_2^n \subseteq \ell_K(A)K$. Assume not, then there exists $\mathbf{x} \in E$ such that $\|\mathbf{x}\|_K = \sup_{\mathbf{y} \in K^*} |\langle \mathbf{y}, \mathbf{x} \rangle| > \ell_K(A)$. Now pick $\mathbf{y} \in K^*$ achieving $|\langle \mathbf{y}, \mathbf{x} \rangle| = \|\mathbf{x}\|_K$. Then we have that

$$\ell_K(A) < |\langle \mathbf{x}, \mathbf{y} \rangle| \leq \sup_{\mathbf{z} \in AB_2^n} |\langle \mathbf{z}, \mathbf{y} \rangle| = \sup_{\mathbf{z} \in B_2^n} |\langle \mathbf{z}, A^t \mathbf{y} \rangle| = \|A^t \mathbf{y}\|_2$$

But now note that

$$\ell_K(A) = \mathbb{E}[\|AX\|_K^2]^{\frac{1}{2}} \geq \mathbb{E}[|\langle \mathbf{y}, AX \rangle|^2]^{\frac{1}{2}} = \|A^t \mathbf{y}\|_2$$

a clear contradiction. Therefore $AB_2^n \subseteq \ell_K(A)K$ as needed. Now applying the same argument on $E^* = A^{-1}B_2^n$ and K^* , we get that $E^* \subseteq \ell_{K^*}(A^{-1})K^*$. From here via duality, we get that

$$\frac{1}{\ell_{K^*}(A^{-1})}K = (\ell_{K^*}(A^{-1})K^*)^* \subseteq (A^{-1}B_2^n)^* = AB_2^n$$

as needed. □

4.4.3 A Deterministic M-Ellipsoid Construction

In this section, we present the algorithm for computing an M-Ellipsoid of an arbitrary convex body in the oracle model.

We first observe that it suffices to give an algorithm for centrally symmetric K . For a general convex body K , we may replace K by the difference body $K - K$ (which is symmetric). An M-Ellipsoid for $K - K$ remains one for K , as the covering estimates changes by at most a $2^{O(n)}$ factor. To see this, note that for any ellipsoid E we have that $N(K, E) \leq N(K - K, E)$ and that

$$N(E, K) \leq N(E, K - K)N(K - K, K) \leq N(E, K - K)2^{O(n)},$$

To prove the last inequality, we use Lemma 2.3.9 which gives

$$\begin{aligned} N(K - K, K) &\leq \sup 6^n \sup_{\mathbf{c} \in \mathbb{R}^n} \frac{\text{vol}_n(K - K)}{\text{vol}_n((K - K) \cap (K + \mathbf{c}))} \\ &= 6^n \frac{\text{vol}_n(K - K)}{\text{vol}_n(K)} \leq 6^n 4^n = 24^n \end{aligned}$$

where the inequality $\text{vol}_n(K - K) \leq 4^n \text{vol}_n(K)$ follows from the classical Rogers-Shephard inequality [111] (see Theorem 2.3.7).

Lastly, given a (\mathbf{a}_0, r, R) -centered convex body K presented by a weak membership oracle O_K , we can construct a weak membership oracle for $K - K$ in polynomial time (see [56]). In the what follows, we now assume that K is symmetric and given by a weak membership oracle.

Our algorithm has two main components: a subroutine to compute an approximate Lewis ellipsoid for a norm given by a convex body, and an implementation of the iteration that makes this ellipsoid converge to an M-Ellipsoid of the original convex body.

4.4.3.1 Approximating the ℓ -norm

Our approximation of the ℓ_K norm is as follows:

$$\tilde{\ell}_K(A) = \sum_{\mathbf{x} \in \{-1, 1\}^n} \frac{1}{2^n} \|A\mathbf{x}\|_K.$$

Note that $\tilde{\ell}_K(A)$ can be deterministically computed using 2^n queries to a weak distance oracle D_K for $\|\cdot\|_K$ using polynomial space (by simply iterating over $\{-1, 1\}^n$). Furthermore, by the guarantees on D_K , for $\epsilon > 0$, we have that

$$\begin{aligned} \left| \sum_{\mathbf{x} \in \{-1, 1\}^n} \frac{1}{2^n} (D_K(A\mathbf{x}, \epsilon) - \|A\mathbf{x}\|_K) \right| &\leq \sum_{\mathbf{x} \in \{-1, 1\}^n} \frac{1}{2^n} |D_K(A\mathbf{x}, \epsilon) - \|A\mathbf{x}\|_K| \\ &\leq \sum_{\mathbf{x} \in \{-1, 1\}^n} \frac{1}{2^n} \epsilon \min\{1, \|A\mathbf{x}\|_K\} = \epsilon \min\{1, \tilde{\ell}_K(A)\} \end{aligned}$$

Hence $\tilde{\ell}_K$ can be estimated to any desired accuracy in $2^n \text{poly}(\cdot)$ time and polynomial space.

The next lemma is essentially folklore, we give a known proof here.

Lemma 4.4.9. *For a symmetric convex body K and any $A \in \mathbb{R}^{n \times n}$, we have*

$$\sqrt{\frac{2}{\pi}} \tilde{\ell}_K(A) \leq \ell_K(A) \leq 4\sqrt{\frac{\pi}{2}} (1 + \log_2 d_{BM}(K, B_2^n)) \tilde{\ell}_K(A).$$

Proof. Let $\mathbf{g}_1, \dots, \mathbf{g}_n$ denote i.i.d. $N(0, 1)$ Gaussians, let $\mathbf{u}_1, \dots, \mathbf{u}_n$ denote i.i.d. uniform $\{-1, 1\}$ random variables and let $A_1, \dots, A_n \in \mathbb{R}^n$ denote the columns of A .

We begin with the lower bound. We relate a classical comparison theorem:

$$\mathbb{E}[f(\mathbf{u}_1, \dots, \mathbf{u}_n)] \leq \mathbb{E}[f(\sqrt{\frac{\pi}{2}} \mathbf{g}_1, \dots, \sqrt{\frac{\pi}{2}} \mathbf{g}_n)] \quad (4.4.3)$$

for any convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Letting $f(\mathbf{x}_1, \dots, \mathbf{x}_n) = \|\sum_i A_i \mathbf{x}_i\|_K$, we see that

$$\tilde{\ell}_K(A) = \mathbb{E}[\|\sum_i A_i \mathbf{u}_i\|_K^2]^{\frac{1}{2}} \leq \mathbb{E}[\|\sqrt{\frac{\pi}{2}} \sum_i A_i \mathbf{g}_i\|_K^2]^{\frac{1}{2}} = \sqrt{\frac{\pi}{2}} \ell_K(A)$$

as needed.

For the upper bound, we see that

$$\begin{aligned} \ell_K(A) &\leq 4(1 + \log_2 d_{BM}(K, B_2^n)) \sup \left\{ \sum_i \langle A_i, \mathbf{y}_i \rangle : \mathbb{E}[\|\sum_i \mathbf{g}_i \mathbf{y}_i\|_{K^*}^2]^{\frac{1}{2}} \leq 1 \right\} \\ &\leq 4\sqrt{\frac{\pi}{2}} (1 + \log_2 d_{BM}(K, B_2^n)) \sup \left\{ \sum_i \langle A_i, \mathbf{y}_i \rangle : \mathbb{E}[\|\sum_i \mathbf{u}_i \mathbf{y}_i\|_{K^*}^2]^{\frac{1}{2}} \leq 1 \right\} \\ &\leq 4\sqrt{\frac{\pi}{2}} (1 + \log_2 d_{BM}(K, B_2^n)) \mathbb{E}[\|\sum_i \mathbf{u}_i A_i\|_K^2]^{\frac{1}{2}} = 4\sqrt{\frac{\pi}{2}} (1 + \log_2 d_{BM}(K, B_2^n)) \tilde{\ell}_K(A) \end{aligned}$$

Here, the first inequality follows by Lemma 4.4.2. The second inequality, follows from the comparison inequality (4.4.3), setting the convex function f to $f(\mathbf{x}_1, \dots, \mathbf{x}_n) = \|\sum_i \mathbf{x}_i \mathbf{y}_i\|_{K^*}$. The last inequality follows from the following weak duality relation:

$$\begin{aligned} \sum_i \langle A_i, \mathbf{y}_i \rangle &= \mathbb{E} \left[\left\langle \sum_i \mathbf{u}_i A_i, \sum_j \mathbf{u}_j \mathbf{y}_j \right\rangle \right] \leq \mathbb{E} \left[\|\sum_i \mathbf{u}_i A_i\|_K \|\sum_j \mathbf{u}_j \mathbf{y}_j\|_{K^*} \right] \\ &\leq \mathbb{E} \left[\|\sum_i \mathbf{u}_i A_i\|_K^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\|\sum_j \mathbf{y}_j \mathbf{u}_j\|_{K^*}^2 \right]^{\frac{1}{2}} \leq \ell_K(A). \end{aligned}$$

□

The next lemma is a strengthening due to Pisier, using Proposition 8 from [104]. While it is not critical for our results, we use this stronger bound in our analysis.

Lemma 4.4.10. *For a symmetric convex body K and any $A \in \mathbb{R}^{n \times n}$, we have*

$$\frac{1}{\sqrt{\frac{\pi}{2}}} \tilde{\ell}_K(A) \leq \ell_K(A) \leq c_1 \tilde{\ell}_K(A) \sqrt{1 + \log d_{BM}(K, B_2^n)}$$

where c_0, c_1 are absolute constants. Furthermore, by duality, we get that

$$\frac{1}{c_1 \sqrt{1 + \log d_{BM}(K, B_2^n)}} \tilde{\ell}_K^*(A) \leq \ell_K^*(A) \leq \sqrt{\frac{\pi}{2}} \tilde{\ell}_K^*(A).$$

4.4.3.2 Convex Program for the ℓ -Ellipsoid

To compute the approximate ℓ -Ellipsoid we use the following convex program:

$$\begin{aligned} \max \det(A)^{\frac{1}{n}} \\ \text{s.t.} \\ A \succeq 0 \\ \tilde{\ell}_K(A) \leq 1 \end{aligned} \tag{4.4.4}$$

Here the main thing we change is that we replace the ℓ -norm with $\tilde{\ell}_K$. This will suffice for our purposes. We optimize over only positive semidefinite matrices (unlike Lewis' program 4.4.2). This enables us to ensure convexity of program while maintaining the desired properties for the optimal solution. For convenience we use $\det(\cdot)^{1/n}$ as the objective function and clearly this makes no essential difference.

4.4.3.3 The Algorithm

Given a convex body K , we put it in approximate John position using the Ellipsoid algorithm in polynomial time [56], so that $B_2^n \subseteq K \subseteq nB_2^n$. We then use the above procedure, which gives an algorithmic implementation of Milman's M-Ellipsoid construction. In the description below, by $\log^{(i)} n$ we mean the i 'th iterated logarithm, i.e., $\log^{(1)} n = 1, \log^{(2)} n = \log \log n$ and so on.

M-Ellipsoid.

- (1) Let $K_1 = K$ and $T = \log^* n$
- (2) For $i = 1 \dots T - 1$,
 - (a) Compute an approximate ℓ -Ellipsoid of K_i using the convex program (4.4.4) to get an approximately optimal transformation A_i (the corresponding ellipsoid is $A_i B_2^n$).
 - (b) Set
$$r_{in} = \frac{\sqrt{n}}{\log^{(i)}(n) \tilde{\ell}_{K_i}(A_i)} \text{ and } r_{out} = \log^{(i)}(n) \frac{\tilde{\ell}_{K_i^*}(A_i^{-1})}{\sqrt{n}}.$$
 - (c) Define
$$K_{i+1} = \text{conv}\{K_i \cap r_{out} A_i B_2^n, r_{in} A_i B_2^n\}.$$
- (3) Output $E = \frac{\sqrt{n}}{\tilde{\ell}_{K_{T-1}}(A_{T-1})} A_{T-1} B_2^n$ as the M-Ellipsoid.

Figure 4.1: The M-Ellipsoid Algorithm

4.4.4 Analysis

We note that the time complexity of the algorithm is bounded by $\text{poly}(n)2^{O(n)}$ and the space complexity is polynomial in n . In fact, the only step that takes exponential time is the evaluation of the ℓ -norm constraint of the SDP. This evaluation happens a polynomial number of times. The rest of computation involves applying the ellipsoid algorithm and computing oracles for successive bodies (for K_{i+1} given an oracle for K_i), both of which are fairly straightforward [56]. In particular, we build an oracle for the intersection of two convex bodies given by oracles and for the convex hull of two convex bodies given by oracles. The oracle for a body consists of a membership test and a bound on the ratio between two balls that sandwich the body. Our analysis below provides sandwiching bounds and the complexity of the oracle grows as $n^{O(i)}$ in the i 'th iteration, for a maximum of $n^{O(\log^* n)} = 2^{o(n)}$.

We begin by showing that Lewis's bound (Theorem 4.4.1) is robust to approximation and works when restricted to positive semi-definite transformations. This

allows us to establish the desired properties for approximate optimizers of the convex program (4.4.4).

Lemma 4.4.11. *Let K be such that $B_2^n \subseteq K \subseteq nB_2^n$ and $A \in \mathbb{R}^{n \times n}$, be a $(1 - \epsilon)$ -approximate optimizer for the convex program (4.4.4), i.e. $\det(A)^{\frac{1}{n}} \geq (1 - \epsilon)OPT$. Then for $\epsilon \leq 1/36n^4$, we have that*

$$\tilde{\ell}_K(A)\tilde{\ell}_K^*(A^{-1}) \leq n(1 + 6n^2\sqrt{\epsilon}) \leq 2n.$$

Proof. For simplicity of notation, we write $\tilde{\ell}_K(T)$ as $\alpha(T)$ for $T \in \mathbb{R}^{n \times n}$. Take $T \in \mathbb{R}^{n \times n}$ (not necessarily positive semidefinite) satisfying $\alpha(T) \leq 1$. Let $\|T\|_F = \sqrt{\sum_{i,j} T_{ij}^2}$ denote the frobenius norm of T , and $\|T\|_2 = \sup_{\mathbf{x} \in B_2^n} \|T\mathbf{x}\|_2$ denote the operator norm of T .

Claim: $\alpha(T) \leq \|T\|_F \leq n\alpha(T)$.

Proof. Let U denote a uniform vector in $\{-1, 1\}^n$. Since $\frac{1}{n}\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_K$ for any $\mathbf{x} \in \mathbb{R}^n$, we have that

$$\alpha(T) = \mathbb{E}[\|UT\|_K^2]^{\frac{1}{2}} \geq \frac{1}{n} \mathbb{E}[\|UT\|_2^2]^{\frac{1}{2}} = \frac{1}{n} \|T\|_F.$$

Now using the inequality $\|\mathbf{x}\|_K \leq \|\mathbf{x}\|_2$ for $\mathbf{x} \in \mathbb{R}^n$, a similar argument yields $\alpha(T) \leq \|T\|_F$. \square

First note that $I_n/\alpha(I_n)$ is a feasible solution to (4.4.4) satisfying

$$\det\left(\frac{I_n}{\alpha(I_n)}\right)^{\frac{1}{n}} = \frac{1}{\alpha(I_n)} \geq \frac{1}{\|I_n\|_F} = \frac{1}{\sqrt{n}}.$$

Let $A_{OPT} \succeq 0$ denote the optimal solution to (4.4.4). Since $\det(A_{OPT}) \geq \frac{1}{\sqrt{n}}$, we clearly have that $A_{OPT} \succ 0$. Therefore for $\delta > 0$ small enough we have that $A_{OPT} + \delta T \succeq 0$. From this, we see that $(A_{OPT} + \delta T)/\alpha(A_{OPT} + \delta T)$ is also feasible for (4.4.4) as $\alpha((A_{OPT} + \delta T)/\alpha(A_{OPT} + \delta T)) = 1$. Since A_{OPT} is the optimal solution, we have that

$$\det\left(\frac{A_{OPT} + \delta T}{\alpha(A_{OPT} + \delta T)}\right)^{\frac{1}{n}} \leq \det(A_{OPT})^{\frac{1}{n}}.$$

Rewriting this and using the triangle inequality,

$$\begin{aligned} \det(A_{OPT} + \delta T)^{\frac{1}{n}} &\leq \det(A_{OPT})^{\frac{1}{n}} \alpha(A_{OPT} + \delta T) \leq \det(A_{OPT})^{\frac{1}{n}} (\alpha(A_{OPT}) + \delta \alpha(T)) \\ &\leq \det(A_{OPT})^{\frac{1}{n}} (1 + \delta). \end{aligned}$$

Dividing by $\det(A_{OPT})^{\frac{1}{n}}$ on both sides, we get that

$$\det(I_n + \delta A_{OPT}^{-1} T)^{\frac{1}{n}} \leq 1 + \delta. \quad (4.4.5)$$

Since both sides are equal at $\delta = 0$, we must have the same inequality for the derivatives with respect to δ at 0. This yields

$$\frac{1}{n} \operatorname{tr}(A_{OPT}^{-1} T) \leq 1 \Leftrightarrow \operatorname{tr}(A_{OPT}^{-1} T) \leq n \quad (4.4.6)$$

Up to this point the proof is essentially the same as Lewis' proof of Theorem 4.4.1.

We now depart from that proof to account for approximately optimal solutions.

Claim: $\|A_{OPT}^{-1}\|_2 \leq n$.

Proof. Let σ denote the largest eigenvalue of A_{OPT}^{-1} and $\mathbf{v} \in \mathbb{R}^n$ be an associated unit eigenvector. Since $A_{OPT} \succ 0$, we have that $A_{OPT}^{-1} \succ 0$, and hence $\sigma = \|A^{-1}\|_2$. Now note that $A_{OPT} + \delta \mathbf{v} \mathbf{v}^T \succ 0$ for any $\delta \geq 0$, and that $\alpha(\mathbf{v} \mathbf{v}^T) \leq \|\mathbf{v} \mathbf{v}^T\|_F = \|\mathbf{v}\|_2^2 = 1$. Therefore by Equation (4.4.6), we have that

$$n \geq \operatorname{tr}(A^{-1}(\mathbf{v} \mathbf{v}^T)) = \operatorname{tr}(\sigma \mathbf{v} \mathbf{v}^T) = \sigma$$

as needed. □

Claim: $A^{-1} \preceq (1 + 6\sqrt{n\epsilon})A_{OPT}^{-1}$.

Proof. Since A is $(1 - \epsilon)$ -approximate maximizer to (4.4.4) we have that

$$\det(A)^{\frac{1}{n}} \geq (1 - \epsilon) \det(A_{OPT})^{\frac{1}{n}} \Rightarrow \det(A) \geq (1 - n\epsilon) \det(A_{OPT})$$

We begin by proving by proving $A \succeq (1 - 3\sqrt{n\epsilon})A_{OPT}$. Now note that

$$A \succeq (1 - 3\sqrt{n\epsilon})A_{OPT} \Leftrightarrow A_{OPT}^{-\frac{1}{2}}AA_{OPT}^{-\frac{1}{2}} \succeq (1 - 3\sqrt{n\epsilon})I_n$$

Hence letting $B = A_{OPT}^{-\frac{1}{2}}AA_{OPT}^{-\frac{1}{2}}$, it suffices to show that $B \succeq (1 - 3\sqrt{n\epsilon})I_n$. From here, we note that $1 \geq \det(B) = \det(A)/\det(A_{OPT}) \geq (1 - n\epsilon)$. Now from Equation (4.4.6), we have that

$$\text{tr}(B) = \text{tr}(A_{OPT}^{-\frac{1}{2}}AA_{OPT}^{-\frac{1}{2}}) = \text{tr}(A_{OPT}^{-1}A) \leq n$$

Let $\sigma_1, \dots, \sigma_n \geq 0$ denote the eigen values of B in non-increasing order. We first note that $\sigma_n \leq 1$ since otherwise

$$\det(B) = \prod_{i=1}^n \sigma_i \geq \sigma_n^n > 1$$

a contradiction. Furthermore, since $B \succ 0$, we have that $0 < \sigma_n \leq 1$. So we may write $\sigma_n = 1 - \epsilon_0$, for $1 > \epsilon_0 \geq 0$. Now since $\sum_{i=1}^n \sigma_i = \text{tr}(B) \leq n$, by the arithmetic mean - geometric mean inequality we have that

$$\det(B) = \sigma_n \prod_{i=1}^{n-1} \sigma_i = (1 - \epsilon_0) \prod_{i=1}^{n-1} \sigma_i \leq (1 - \epsilon_0) \left(\frac{\sum_{i=1}^{n-1} \sigma_i}{n-1} \right)^{n-1} \leq (1 - \epsilon_0) \left(1 + \frac{\epsilon_0}{n-1} \right)^{n-1}$$

Using the inequality $1 + x \leq e^x \leq 1 + x + \frac{e-1}{2}x^2$ for $x \in [-1, 1]$, we get that

$$\begin{aligned} (1 - \epsilon_0) \left(1 + \frac{\epsilon_0}{n-1} \right)^{n-1} &\leq (1 - \epsilon_0) e^{\epsilon_0} \leq (1 - \epsilon_0) \left(1 + \epsilon_0 + \frac{e-1}{2} \epsilon_0^2 \right) \\ &= 1 - \frac{3-e}{2} \epsilon_0^2 - \frac{e-1}{2} \epsilon_0^3 \leq 1 - \frac{3-e}{2} \epsilon_0^2 \end{aligned}$$

From this we get that

$$1 - \frac{3-e}{2} \epsilon_0^2 \geq \det(B) \geq (1 - n\epsilon) \Rightarrow \epsilon_0 \leq \sqrt{\frac{2}{3-e} n\epsilon} \leq 3\sqrt{n\epsilon}$$

Therefore $\sigma_n = 1 - \epsilon_0 \geq 1 - 3\sqrt{n\epsilon} \Rightarrow B \succeq (1 - 3\sqrt{n\epsilon})I_n \Rightarrow A \succeq (1 - 3\sqrt{n\epsilon})A_{OPT}$ as needed. From here we get that

$$A^{-1} \preceq \left(\frac{1}{1 - 3\sqrt{n\epsilon}} \right) A_{OPT}^{-1} \preceq (1 + 6\sqrt{n\epsilon})A_{OPT}^{-1}$$

for $\epsilon \leq 1/36n$, proving the claim. \square

Now take $T \in \mathbb{R}^{n \times n}$ satisfying $\alpha(T) \leq 1$. By the first claim, we note that $\|T\|_F \leq n\alpha(T) \leq n$. Now by Equation (4.4.6), we have that

$$\mathrm{tr}(A^{-1}T) = \mathrm{tr}(A_{OPT}^{-1}T) + \mathrm{tr}((A^{-1} - A_{OPT}^{-1})T) \leq n + \|A^{-1} - A_{OPT}^{-1}\|_F \|T\|_F \leq n + n\|A^{-1} - A_{OPT}^{-1}\|_F$$

We bound the second term using the previous claim. Since $A^{-1} \preceq (1 + 6\sqrt{n\epsilon})A_{OPT}^{-1}$, we have that $A^{-1} - A_{OPT}^{-1} \preceq 6\sqrt{n\epsilon}A_{OPT}^{-1}$, and hence

$$\|A^{-1} - A_{OPT}^{-1}\|_F \leq \sqrt{n}\|A^{-1} - A_{OPT}^{-1}\|_2 \leq 6n\sqrt{\epsilon}\|A_{OPT}^{-1}\|_2 \leq 6n^2\sqrt{\epsilon}$$

Using this bound, we get

$$\mathrm{tr}(A^{-1}T) \leq n + 6n^3\sqrt{\epsilon} = n(1 + 6n^2\sqrt{\epsilon})$$

for any $T \in \mathbb{R}^{n \times n}$ satisfying $\alpha(T) \leq 1$. Thus we get that $\alpha^*(A^{-1}) \leq n(1 + 6n^2\sqrt{\epsilon})$.

Together with the constraint $\alpha(A) \leq 1$, the conclusion of the lemma follows. \square

Theorem 4.4.12. *Let A be a $(1 - \epsilon)$ -approximate optimizer to the convex program (4.4.4) for $\epsilon \leq 1/(36n^4)$. Then*

$$\ell_K(A)\ell_{K^*}(A^{-1}) \leq Cn \log^{\frac{3}{2}} d_{BM}(K, B_2^n).$$

for an absolute constant $C > 0$.

Proof. Using Lemma 4.4.11, we have that

$$\tilde{\ell}_K(A)\tilde{\ell}_K^*(A^{-1}) \leq 2n.$$

Next we use the approximation property (Lemma 4.4.10) of $\tilde{\ell}_K$ to derive that

$$\ell_K(A)\ell_{K^*}^*(A^{-1}) \leq Cn\sqrt{\log d_{BM}(K, B_2^n)}.$$

Finally, noting that $A^{-T} = A^{-1}$ (by symmetry of A), we apply Lemma 4.4.2 to infer that

$$\ell_{K^*}(A^{-1}) \leq C\ell_K^*(A^{-1}) \log d_{BM}(K, B_2^n),$$

which completes the proof. \square

Next we turn to proving that the algorithm produces an M-Ellipsoid. While the analysis follows the existence proof to a large extent, we need to handle the various approximations incurred.

To aid in the analysis of Algorithm 4.1 on input $K \subseteq \mathbb{R}^n$, we make some additional definitions. Let $a_i = \log^{(i)} n$ and $T = \log^* n$. Let K_1, \dots, K_T and A_1, \dots, A_T denote the sequence of bodies and transformations generated by the algorithm. Set $K_1^{out} = K_1^{in} = K$, and for $1 \leq i \leq T - 1$ define

$$K_{i+1}^{in} = \text{conv}\{K_i^{in}, r_{in}^i A_i B_2^n\} \quad K_{i+1}^{out} = K_i^{out} \cap r_{out}^i A_i B_2^n$$

where r_{in}^i, r_{out}^i are defined as r_{in}, r_{out} in the i 'th iteration of the main loop in Algorithm 4.1.

By construction, we have the relations

$$K \subseteq K_1^{in} \subseteq \dots \subseteq K_T^{in}, \quad K \supseteq K_1^{out} \supseteq \dots \supseteq K_T^{out}, \quad K_i^{out} \subseteq K_i \subseteq K_i^{in} \quad \forall i \in [T]$$

The proof of the main theorem will be based on the following inductive lemmas which quantify the properties of the sequences of bodies defined above.

Lemma 4.4.13. $\forall i \in [T]$, we have that $d_{BM}(K_i, B_2^n) \leq C(\log^{(i-1)} n)^{\frac{7}{2}}$.

Proof. For the base case, we have that $d_{BM}(K_1, B_2^n) \leq \sqrt{n} \leq Cn^{\frac{7}{2}}$ for any constant $C \geq 1$.

For the general case, by construction of K_{i+1} we have that

$$r_{in}^i A_i B_2^n \subseteq K_{i+1} \subseteq r_{out}^i A_i B_2^n.$$

Therefore,

$$\begin{aligned} d_{BM}(K_{i+1}, B_2^n) &\leq r_{out}^i / r_{in}^i \\ &= a_i^2 \tilde{\ell}_{K_i^*}(A_i^{-1}) \tilde{\ell}_{K_i}(A_i) / n \\ &\leq C_1 a_i^2 \ell_{K_i^*}(A^{-1}) \ell_{K_i}(A_i) / n \quad (\text{by Lemma 4.4.10}) \\ &\leq C_1 (\log^{(i)} n)^2 (\log d_{BM}(K_i, B_2^n))^{\frac{3}{2}}. \quad (\text{by Lemma 4.4.12}) \end{aligned}$$

Using the fact that $\log^{(i)}(n) \geq 1$, $\forall i \in [T-1]$, a direct computation shows that the above recurrence equation implies the existence of a constant $C > 1$ (depending only on C_1) such that the stated bound on $d_{BM}(K_{i+1}, B_2^n)$ holds. \square

Lemma 4.4.14. *For $i \in [T-1]$, we have that*

$$\max \left\{ \frac{\text{vol}(K_i^{\text{out}})}{\text{vol}(K_{i+1}^{\text{out}})}, \frac{\text{vol}(K_{i+1}^{\text{in}})}{\text{vol}(K_i^{\text{in}})} \right\} \leq e^{Cn/\log^{(i)} n}$$

Proof. By Lemma 4.4.4, the fact that $K_i^{\text{out}} \subseteq K_i$, Lemma 4.4.6, Lemma 4.4.10 and Lemma 4.4.13, we have that

$$\begin{aligned} \frac{\text{vol}(K_i^{\text{out}})}{\text{vol}(K_{i+1}^{\text{out}})} &\leq N(K_i^{\text{out}}, r_{\text{out}}^i A_i B_2^n) \leq N(K_i, r_{\text{out}}^i A_i B_2^n) \\ &\leq e^{C(\ell_{K_i^*}(A_i^{-1})/r_{\text{out}}^i)^2} = e^{Cn\ell_{K_i^*}(A_i^{-1})^2/(a_i\tilde{\ell}_{K_i^*}(A_i^{-1}))^2} \\ &\leq e^{Cn\log(d_{BM}(K_i^*, B_2^n))/a_i^2} \leq e^{Cn/\log^{(i)} n} \end{aligned}$$

By Lemma 4.4.8, 4.4.10 and 4.4.13, we see that

$$r_{\text{in}}^i A_i B_2^n \subseteq r_{\text{in}}^i \ell_{K_i^{\text{in}}}(A_i) K_i^{\text{in}} \subseteq r_{\text{in}}^i \ell_{K_i}(A_i) K_i^{\text{in}} \subseteq C_1 \sqrt{n} K_i^{\text{in}}.$$

Next by Lemma 4.4.5, the fact that $K_i \subseteq K_i^{\text{in}}$, Lemma 4.4.7, Lemma 4.4.10 and Lemma 4.4.13, we have that

$$\begin{aligned} \frac{\text{vol}(K_{i+1}^{\text{in}})}{\text{vol}(K_i^{\text{in}})} &\leq C_1 4n^{\frac{3}{2}} N(r_{\text{in}}^i A_i B_2^n, K_i^{\text{in}}) \leq C_1 n^{\frac{3}{2}} N(r_{\text{in}}^i A_i B_2^n, K_i) \\ &\leq C_1 n^{\frac{3}{2}} e^{C(\ell_{K_i}(A_i)r_{\text{in}}^i)^2} = C_1 n^{\frac{3}{2}} e^{Cn\ell_{K_i}(A_i)^2/(a_i\tilde{\ell}_{K_i}(A_i))^2} \\ &\leq C_1 n^{\frac{3}{2}} e^{Cn\log(d_{BM}(K_i, B_2^n))/a_i^2} \leq C_1 n^{\frac{3}{2}} e^{Cn(1/\log^{(i)} n)} \leq e^{Cn/\log^{(i)} n} \end{aligned}$$

\square

We are now ready to complete the proof.

Proof. (of Theorem 4.1.3.) By construction of K_T , we note that

$$r_{\text{in}}^{T-1} A_{T-1} B_2^n \subseteq K_T \subseteq r_{\text{out}}^{T-1} A_{T-1} B_2^n$$

where by Lemma 4.4.13 we have that $r_{out}^{T-1}/r_{in}^{T-1} = O(1)$. Therefore the returned ellipsoid $E = \frac{\sqrt{n}}{\tilde{\ell}_{K_{T-1}(A_{T-1})}} A_{T-1} B_2^n$ (last line of Algorithm 4.1) satisfies that

$$\frac{1}{C}E \subseteq K_T \subseteq CE$$

for an absolute constant $C \geq 1$. Next by Lemma 4.4.4, we have that

$$N(K, E), N(E, K) \leq 3^n \frac{\max\{\text{vol}(K), \text{vol}(E)\}}{\text{vol}(K \cap E)}$$

Now we see that

$$K \subseteq K_T^{in} \subseteq CK_T^{in} \quad E \subseteq CK_T \subseteq CK_T^{in},$$

and that

$$K \supseteq \frac{1}{C}K_T^{out} \quad E \supseteq \frac{1}{C}K_T \supseteq \frac{1}{C}K_T^{out}.$$

Therefore,

$$\frac{\max\{\text{vol}(K), \text{vol}(E)\}}{\text{vol}(K \cap E)} \leq C^{2n} \frac{\text{vol}(K_T^{in})}{\text{vol}(K_T^{out})}.$$

Finally, by Lemma 4.4.14 we have that

$$\frac{\text{vol}(K_T^{in})}{\text{vol}(K_T^{out})} = \prod_{i=1}^{T-1} \frac{\text{vol}(K_{i+1}^{in})}{\text{vol}(K_i^{in})} \frac{\text{vol}(K_i^{out})}{\text{vol}(K_{i+1}^{out})} \leq \prod_{i=1}^{T-1} e^{2Cn/\log^{(i)} n} = 2^{O(n)}.$$

Combining the above inequalities yields the desired guarantee on the algorithm. The time complexity is $2^{O(n)}$, dominated by the time to evaluate the $\tilde{\ell}_K$ -norm. The space is polynomial since all we need to maintain are efficient oracles for the successive bodies K_i , which can be done space-efficiently for the operations of intersection and convex hull used in the algorithm [56]. \square

4.5 An Asymptotically Optimal Volume Algorithm

As noted in the introduction, the result of Theorem 4.1.5, a deterministic $2^{O(n)}$ -approximation for volume, follows directly from Theorem 4.1.3. In this section, we show how to modify our M-Ellipsoid algorithm (based on Milman's iteration) to match this lower bound algorithmically.

In the M-Ellipsoid algorithm of the previous section, we construct a series of convex bodies $K_0 = K, K_1, \dots, K_T$ such that the covering numbers $N(K, K_T)$ and $N(K_T, K)$ are bounded by $2^{O(n)}$ and the final body K_T has $d_{BM}(K_T, B_2^n) < C$ for some constant C . Our modification will construct a similar sequence of bodies, but rather than covering numbers, we will ensure that

$$e^{-C\epsilon n} \text{vol}(K) \leq \text{vol}(K_T) \leq e^{C\epsilon n} \text{vol}(K)$$

and

$$d_{BM}(K_T, B_2^n) \leq C \frac{\ln(1/\epsilon)^{\frac{5}{2}}}{\epsilon^2}.$$

Then we approximate the volume of K_T by finding an approximate ℓ -Ellipsoid E for it, and covering it with translations of a maximal parallelepiped that fits in ϵE . Since this covering will consist of disjoint parallelepipeds, and their union will be contained in $K_T + \epsilon E \subseteq (1 + \epsilon)K_T$, we get the desired approximation. Here is the precise algorithm.

Proof of Theorem 4.1.7. Let $a_i = \log^{(i)} n$. As in Lemma 4.4.13, we bound the Banach Mazur via the following recurrence

$$d_{BM}(K_{i+1}, B_2^n) \leq r_{out}^i / r_{in}^i \leq C \frac{\ln(1/\epsilon)}{\epsilon^2} (\log^{(i)}(n))^2 (\log d_{BM}(K_i, B_2^n))^{\frac{3}{2}}.$$

From the above recurrence a direct computation reveals that for $\forall i \in [T]$,

$$d_{BM}(K_i, B_2^n) \leq C \frac{\ln(1/\epsilon)^{5/2}}{\epsilon^2} (\log^{(i-1)}(n))^{\frac{7}{2}}$$

We now show that the volumes of the K_i bodies changes very slowly. This will enable us to conclude that the volume of K_T is very close to the volume of K .

By Lemmas 4.4.8, 4.4.10 and the above bound on $d_{BM}(K_i, B_2^n)$, we have that

$$r_{in}^i A_i B_2^n \subseteq r_{in}^i \ell_{K_i}(A_i) K_i \subseteq C \frac{\epsilon \sqrt{n \log d_{BM}(K_i, B_2^n)}}{\sqrt{\ln(1/\epsilon) \log^{(i)}(n)}} K_i \subseteq C \epsilon \sqrt{n} K_i$$

and that

$$r_{out}^i A_i B_2^n = C \frac{\sqrt{\ln(1/\epsilon) \log^{(i)}(n)} \tilde{\ell}_{K^*}(A^{-1})}{\epsilon \sqrt{n}} A_i B_2^n \supseteq C \frac{\ell_{K^*}(A^{-1})}{\epsilon \sqrt{n}} A_i B_2^n \supseteq C \frac{1}{\epsilon \sqrt{n}} K_i.$$

Volume(K, ϵ).

- (1) Let $K_1 = K$ and $T = \log^* n$
- (2) For $i = 1 \dots T - 1$,
 - (a) Compute an approximate ℓ -Ellipsoid of K_i using the convex program (4.4.4) to get an approximately optimal transformation A_i (the corresponding ellipsoid is $A_i B_2^n$).
 - (b) Set
$$r_{in} = \frac{\epsilon \sqrt{n}}{\sqrt{\ln(1/\epsilon)} C \log^{(i)}(n) \tilde{\ell}_{K_i}(A_i)} \text{ and } r_{out} = \frac{C \sqrt{\ln(1/\epsilon)} \log^{(i)}(n) \tilde{\ell}_{K_i^*}(A_i^*)}{\epsilon \sqrt{n}}.$$
 - (c) Define
$$K_{i+1} = \text{conv}\{K_i \cap r_{out} A_i B_2^n, r_{in} A_i B_2^n\}.$$
- (3) Compute the ellipsoid $E = r_{in} A_{T-1} B_2^n$ and a maximum volume paralleliped P inscribed in E .
- (4) Call Parallelepiped-Tiling($K_T, \frac{\epsilon}{2} P, 1$) to tile K_T with $\frac{\epsilon}{2} P$.
Output $k \text{vol}(P)$, where k is the computed size of the tiling.

Figure 4.2: Deterministic Volume Algorithm

Therefore if $\epsilon \leq \sqrt{n}/C$, then $K_{i+1} = \text{conv}\{r_{in}^i A_i B_2^n, K_i \cap r_{out}^i A_i B_2^n\} = K_i$. Since this holds for all $i \in [T - 1]$, we get that $K_T = K$ and hence $\text{vol}(K_T) = \text{vol}(K)$.

Now assume that $\epsilon \geq \sqrt{n}/C$. Then for $i \in [T - 1]$, using Lemmas 4.4.4 and 4.4.6, we have,

$$\begin{aligned}
\text{vol}(K_{i+1}) &\geq \text{vol}(K_i \cap r_{out} B_2^n) \\
&\geq \frac{\text{vol}(K_i)}{N(K_i, r_{out}^i B_2^n)} \\
&\geq e^{-C(\ell_{K_i^*}(A_i^{-1})/r_{out}^i)^2} \text{vol}(K_i) \\
&\geq e^{-C(\epsilon^2/\ln(1/\epsilon))n \log d_{BM}(K_i, B_2^n)/a_i^2} \text{vol}(K_i) \\
&\geq e^{-Cn\epsilon/\log^{(i)}(n)} \text{vol}(K_i).
\end{aligned}$$

From the above, we get that

$$\frac{\text{vol}(K_T)}{\text{vol}(K)} = \prod_{i=1}^{T-1} \frac{\text{vol}(K_{i+1})}{\text{vol}(K_i)} \geq \prod_{i=1}^{T-1} e^{-Cn\epsilon/\log^{(i)}(n)} \geq e^{-Cn\epsilon}$$

Next via Lemma 4.4.5, the above containment, and Lemma 4.4.6, we have,

$$\begin{aligned} \text{vol}(K_{i+1}) &\leq \text{vol}(\text{conv}\{K_i, r_{in}B_2^n\}) \\ &\leq C(\epsilon\sqrt{n})nN(r_{in}B_2^n, K_i)\text{vol}(K_i) \\ &\leq C(\epsilon n^{\frac{3}{2}})e^{C(r_{in}^i \ell_K(A_i))^2}\text{vol}(K_i) \\ &\leq C(\epsilon n^{\frac{3}{2}})e^{C(\epsilon^2/\ln(1/\epsilon))n \log d_{BM}(K_i, B_2^n)/a_i^2}\text{vol}(K_i) \\ &\leq C(\epsilon n^{\frac{3}{2}})e^{Cn\epsilon/\log^{(i)}(n)}\text{vol}(K_i). \end{aligned}$$

From this, we get that

$$\frac{\text{vol}(K_T)}{\text{vol}(K)} = \prod_{i=1}^{T-1} \frac{\text{vol}(K_{i+1})}{\text{vol}(K_i)} \leq (C\epsilon n^{\frac{3}{2}})^{\log^*(n)} \prod_{i=1}^{T-1} e^{Cn\epsilon/\log^{(i)}(n)} \leq e^{Cn\epsilon},$$

where the above holds as long as $\epsilon = \Omega(\frac{\log n \log^* n}{n})$ (which we have by assumption).

Combining the above inequalities we get

$$e^{-C\epsilon n}\text{vol}(K) \leq \text{vol}(K_T) \leq e^{C\epsilon n}\text{vol}(K).$$

Let $E = E(A)$ denote the final ellipsoid computed by the algorithm, and let $P = A^{-\frac{1}{2}}[-1, 1]^n$ denote a maximum volume inscribed parallelepiped of E . By construction of E and K_T , we have that $E \subseteq K_T \subseteq C\frac{\ln(1/\epsilon)^{5/2}}{\epsilon^2}E$. By the guarantees on Algorithm Parallelepiped-Tiling (from section 4.2) on input K_T , $\frac{\epsilon}{2}P$, and 1, the outputted tiling is contained in $K_T + (1 + 1)\frac{\epsilon}{2}P \subseteq K_T + \epsilon E \subseteq (1 + \epsilon)K_T$. Hence the estimate outputted by the algorithm lies between $\text{vol}(K_T)$ and $\text{vol}((1 + \epsilon)K_T) = (1 + \epsilon)^n \text{vol}(K_T)$. Thus the overall approximation factor is bounded by $e^{Cn\epsilon}$ as desired (setting $\epsilon = c_1\epsilon'$, for $c_1 < 1$ sufficiently small, yields a $(1 + \epsilon')^n$ approximation).

Finally, the running time of the algorithm is dominated by the time to compute the covering. Noting that $\text{vol}(E) = (\sqrt{\frac{\pi e}{2}}(1 + o(1)))^n \text{vol}(P) = 2^{O(n)} \text{vol}(P)$, the size

of the covering is bounded by

$$\begin{aligned} \frac{\text{vol}(K_T + \epsilon P)}{\text{vol}(\frac{\epsilon}{2}P)} &\leq 2^n(1 + \epsilon)^n \frac{\text{vol}(K_T)}{\text{vol}(P)} \leq C^n(1 + \epsilon)^n \frac{\text{vol}(K_T)}{\text{vol}(E)} \\ &\leq C^n(1 + \epsilon)^n (\ln(1/\epsilon))^{5/2} / \epsilon^2)^n = (1/\epsilon)^{O(n)}. \end{aligned}$$

Hence by the guarantees on algorithm Parallelepiped tiling, the time to compute the covering is $(1/\epsilon)^{O(n)}$. Lastly, the space needed to compute the size of the covering is polynomial, as we only need to keep track of the count as we iterate over elements of the covering. \square

4.6 Computing an Approximate Center of Mass

In this section, we give a Las Vegas algorithm which finds an approximate center of mass for any convex body. More precisely, for a convex body $K \subseteq \mathbb{R}^n$, we wish to output a point $\mathbf{b} \in K$ such that $\text{vol}_n(K) \leq 2^{O(n)} \text{vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K))$. We remember that $\text{vol}_n(K) \leq 2^n \text{vol}_n((K - \mathbf{b}(K)) \cap (\mathbf{b}(K) - K))$ (Theorem 2.3.7), where $\mathbf{b}(K)$ is the centroid of K . Given this, we can apply random sampling techniques to estimate the centroid, and get an approximate center of mass with high probability. In future chapters, our lattice algorithms will make critical use of approximate centers of mass. To avoid byzantine failures we will prefer a Las Vegas algorithm for computing such centers. Our strategy here will be to use the randomized sampling methods to approximate the centroid and certify it using the volume estimation algorithm from the previous section.

We give the algorithm below.

Algorithm 4.6 Approx-Mass-Center(K)

Input: A weak membership oracle O_K for a (\mathbf{a}_0, r, R) -centered convex body $K \subseteq \mathbb{R}^n$.

Output: \mathbf{b} satisfying the conditions of Theorem 4.6.1.

1: $\mathbf{b} \leftarrow \text{Estimate-Centroid}(K, \frac{1}{n})$.

 If Estimate-Centroid returns FAIL, restart; else, continue.

2: $V_1 \leftarrow \text{Volume}(K - K, \frac{1}{10})$; $V_2 \leftarrow \text{Volume}((K - \mathbf{b}) \cap (\mathbf{b} - K), \frac{1}{10})$.

3: If $V_1/2^n > (4.5)^n V_2$, restart; else, **return** \mathbf{b} .

Theorem 4.6.1 (Correctness of Approx-Mass-Center). *Given an (\mathbf{a}_0, r, R) -centered convex body $K \subseteq \mathbb{R}^n$, in expected $2^{O(n)}$ $\text{poly}(\cdot)$ time, using $\text{poly}(\cdot)$ space and randomness, algorithm 4.6 outputs a vector $\mathbf{b} \in K$ satisfying (for n large enough)*

$$(1) \text{ vol}_n(K) \leq 5^n \text{ vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K)).$$

$$(2) K \text{ is } (\mathbf{b}, \frac{r}{2(n+1)\sqrt{n}}, 2R) \text{ centered.}$$

Analysis of Approx-Mass-Center.

Correctness: We show that if the algorithm successfully returns \mathbf{b} , then \mathbf{b} satisfies both (1) and (2). We first note that (2) follows directly from the guarantees on the algorithm Estimate-Centroid. Next, note that by the guarantees on algorithm Volume we have that

$$V_1 \leq \text{vol}_n(K - K) \leq (1 + \epsilon)^n V_1 \quad \text{and} \quad V_2 \leq \text{vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K)) \leq (1 + \epsilon)^n V_2 \quad (4.6.1)$$

for $\epsilon = \frac{1}{10}$. Next, by the Rogers-Shephard inequality and the Brunn-Minkowski inequality, we have that

$$2^n \text{vol}_n(K) \leq \text{vol}_n(K - K) \geq 4^n \text{vol}_n(K). \quad (4.6.2)$$

Now assuming the final test passes, we have that

$$\begin{aligned} V_1/2^n \leq (4.5)^n V_2 &\Rightarrow \text{vol}_n(K)/(1 + \epsilon)^n \leq (4.5)^n \text{vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K)) \\ &\Rightarrow \text{vol}_n(K) \leq 5^n \text{vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K)). \end{aligned}$$

Therefore the output of the algorithm is correct, as needed.

Runtime: For the runtime, we see that the calls to Estimate-Centroid run in $\text{poly}(\cdot)$ time and space and the calls to algorithm Volume run in $2^{O(n)}$ $\text{poly}(\cdot)$ time and $\text{poly}(\cdot)$ space. Therefore, to achieve the desired runtime, it suffices to show that the algorithm's loop executes at most $O(1)$ times.

To prove this, we will show that if Estimate-Centroid outputs \mathbf{b} satisfying $\mathbf{b} \in \mathbf{b}(K) + \frac{1}{n+1}E_K$, then the loop tests succeed. By the guarantees on Estimate-Centroid on input K and error probability $\frac{1}{n}$, this event happens with probability at least $1 - \frac{1}{n}$. This will therefore suffice to show that the number of loop iterations is $O(1)$ on expectation.

Since $\mathbf{b} \in \mathbf{b}(K) + \frac{1}{n+1}E_K$, by theorem 4.3.6 we have that

$$\text{vol}_n((\mathbf{b} - K) \cap (\mathbf{b} - K)) \geq \left(1 - \frac{1}{n+1}\right)^n 2^{-n} \text{vol}_n(K) \geq \frac{1}{e} 2^{-n} \text{vol}_n(K).$$

Therefore since $(1+\epsilon)^n V_2 \geq \text{vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K))$ and that $2^n \text{vol}_n(K) \geq 2^{-n} \text{vol}_n(K - K) \geq 2^{-n} V_1$, we have that

$$e 4^n (1+\epsilon)^n V_2 \geq e 4^n \text{vol}_n((K - \mathbf{b}) \cap (\mathbf{b} - K)) \geq 2^n \text{vol}_n(K) \geq 2^{-n} \text{vol}_n(K - K) \geq 2^{-n} V_1$$

Next, we note that

$$e 4^n (1+\epsilon)^n = e 4^n \left(1 + \frac{1}{10}\right)^n = e 4.4^n \leq 4.5^n$$

for $n \geq 45$. Therefore the test at the end of the loop passes with probability at least $1 - \frac{1}{n}$, for n large enough, as needed. We note that the only randomness in the algorithm is used by Estimate-Centroid, and since Estimate-Centroid runs in polynomial time, the amount of randomness used in the algorithm is polynomial as needed. \square

Acknowledgments. We would like to express our gratitude to Gideon Schechtman for suggesting the use of parallelepiped tilings to build the M-Ellipsoid covering, to Matthias Köppe for suggesting the use of reverse search to save space in our covering algorithms, to Boaz Klartag for informing us that his M-Ellipsoid construction could be used to build the M-Ellipsoid algorithmically, to Grigoris Paouris for informing us of inequality 4.4.10 and to Assaf Naor for helping us locate an appropriate reference for it.

4.7 Conclusion

The study of different types of ellipsoids for convex bodies and the complexity of computing them has played a fundamental role in both convex geometry and computer science. The complexity of computing John and inertial ellipsoids for example, has been well studied. In this Chapter, we have focused our attention on computing another fundamental ellipsoid in convex geometry: the M-Ellipsoid. To the best of our knowledge, this problem had not been studied previously.

Here, we have given two different algorithms for computing M-Ellipsoids: the first being a randomized polynomial time algorithm, based on a construction of Klartag [77], and the second being a deterministic $2^{O(n)}$ time algorithm, based on Milman's original construction [92]. A main motivation for studying M-Ellipsoid constructions has been to gain a better understanding of what properties of a convex body are deterministically computable. Our main result from this perspective, has been to show that the M-Ellipsoid (in particular its deterministic construction), can be used to give a nearly optimal algorithm for estimating the volume of a symmetric convex body. Lastly, as a crucial tool, we have developed an algorithm to efficiently cover any convex body by an ellipsoid. This algorithm will play a central role in the applications to lattice problems and integer programming in the next chapters.

Future Research. A first future research direction is to find more optimized algorithms than the ones given here. We have made very little effort in this Chapter to optimize the quality of the outputted M-Ellipsoid, and understanding how small the product of covering estimates $N(K, E)N(E, K)$ can be made in an algorithmic and deterministic manner is a very interesting question. Following this, to deterministically estimate the volume of a symmetric convex body to within $(1 + \epsilon)^n$ our algorithm requires roughly $(1 + 1/\epsilon)^{2n}$ time (modulo polylogarithmic factors in ϵ). From the Bárány-Füredi lower bound [52], the optimal dependence on ϵ should be

$(1 + \epsilon)^{\frac{n}{2}}$. Closing this gap would seem to require stronger methods than the ones described here, and is an interesting open problem.

A next series of questions, is on the subject of generalizing the algorithms given here. Firstly, the volume algorithm presented here only works for symmetric convex bodies. A first question is therefore whether the same can be done for asymmetric convex bodies in the same time and space complexity. Second, our current covering algorithm only works for covering convex bodies by ellipsoids (or more precisely, cuboids). Given two convex bodies K_1, K_2 , is there a general method for covering K_1 by K_2 in a near optimal fashion (i.e. of size roughly $N(K_1, K_2)$)?

Lastly, we would like explore algorithmic questions on convex bodies beyond the oracle model. In particular, given the work presented here, a tantalizing question is whether one can compute M-Ellipsoids for explicit polytopes (i.e. whose descriptions are given explicitly as input) in polynomial time. The main complexity of Milman's construction for example, is computing the expectation

$$\mathbb{E}[\|TU\|_K] = \sum_{\mathbf{x} \in \{-1,1\}^n} \frac{1}{2^n} \|T\mathbf{x}\|_K$$

If K is a symmetric polytope, we can express $K = \{\mathbf{z} : |A\mathbf{z}| \leq 1\}$. We note that $\|\mathbf{x}\|_K = \|A\mathbf{x}\|_\infty$, and hence $\mathbb{E}[\|TU\|_K] = \mathbb{E}[\|ATU\|_\infty]$. Here it seems very likely that one may be able to shortcut the complexity of computing this expectation by directly inspecting the matrix AT , avoiding the cost associated with directly evaluating over all of $\{-1, 1\}^n$. Furthermore, we note that for symmetric bodies, the problem computing an M-Ellipsoid is self dual, i.e. it is equivalent to compute an M-Ellipsoid for either K or K^* . Therefore, from the perspective of symmetric polytopes, the complexity of the problem is the same whether the polytope is given in either convex hull or inequality representation.

CHAPTER V

EFFICIENT DETERMINISTIC ALGORITHMS FOR LATTICE PROBLEMS

We give a new algorithm for enumerating lattice points in any convex body, and give applications to several classic lattice problems, including the Shortest and Closest Vector Problems (SVP and CVP, respectively). Our enumeration technique relies on the classical *M-Ellipsoid* concept from convex geometry (presented in the previous chapter), and uses as a crucial subroutine the recent algorithm of Micciancio and Voulgaris [91] for lattice problems in the ℓ_2 norm. Additionally, we give a novel algorithm to “sparsify” an input lattice at any desired distance while approximately maintaining the lattice’s metric structure.

As applications, we give deterministic single exponential time and space algorithms for solving exact SVP, exact CVP when the target is “close”, and $(1 + \epsilon)$ -CVP, on n -dimensional lattices *in any norm*. Our approach yields the first deterministic alternative to the “AKS Sieve” (Ajtai, Kumar and Sivakumar) for exact SVP and $(1 + \epsilon)$ -CVP [2, 3] for norms other than ℓ_2 .

This Chapter is based on work from the paper [36] (joint with Santosh Vempala, Chris Peikert) as well as subsequent extensions.

5.1 Introduction

The Shortest and Closest Vector Problems (SVP and CVP, respectively) on lattices are central algorithmic problems in the geometry of numbers, with applications to Integer Programming [84], factoring polynomials over the rationals [83], cryptanalysis

(e.g., [101, 66, 99]), and much more. (See section 2.4 for appropriate lattice definitions) The SVP is simply: given a lattice \mathcal{L} represented by a basis, find a nonzero $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\|$ is minimized, where $\|\cdot\|$ denotes a particular norm on \mathbb{R}^n . The CVP is an inhomogeneous analogue of SVP: given a lattice \mathcal{L} and a point $\mathbf{t} \in \mathbb{R}^n$, find some $\mathbf{v} \in \mathcal{L}$ that minimizes $\|\mathbf{v} - \mathbf{t}\|$. In these problems, one often uses the Euclidean (ℓ_2) norm, but many applications require other norms like ℓ_p or, most generally, the norm defined by a convex body $K \ni 0$ as $\|\mathbf{x}\|_K = \inf\{s \geq 0 : \mathbf{x} \in sK\}$. Indeed, general norms arise quite often in the study of lattices; for example, the “flatness theorem” in Integer Programming — which states that every lattice-free convex body has lattice width bounded by a function of the dimension alone — is a statement about SVP in general norms (see Chapter 6).

Much is known about the computational complexity of SVP and CVP, in both their exact and approximation versions. On the negative side, SVP is NP-hard (in ℓ_2 , under randomized reductions) to solve exactly, or even to approximate to within any constant factor [1, 23, 89, 75]. Many more hardness results are known for other ℓ_p norms and under stronger complexity assumptions than $P \neq NP$ (see, e.g., [125, 40, 109, 62]). CVP is NP-hard to approximate to within $n^{c/\log \log n}$ factors for some constant $c > 0$ [4, 41, 40], where n is the dimension of the lattice. Therefore, we do not expect to solve (or even closely approximate) these problems efficiently in high dimensions. Still, algorithms providing weak approximations or having super-polynomial running times are the foundations for the many applications mentioned above.

The celebrated LLL algorithm [83] and variants [118] give $2^{n/\text{polylog}(n)}$ approximations to SVP and CVP in ℓ_2 , in $\text{poly}(n)$ time. For exact SVP and CVP in the ℓ_2 norm, Kannan’s algorithm [70] gives a solution in deterministic $2^{O(n \log n)}$ time and $\text{poly}(n)$ space. This performance remained essentially unchallenged until the breakthrough randomized “sieve” algorithm of Ajtai, Kumar, and Sivakumar [2], which provides a

$2^{O(n)}$ -time and -space solution for exact SVP; moreover, the algorithm was shown, in a sequence of works, to generalize to essentially any norm [18, 5, 33]. For CVP, in roughly the same sequence [3, 18, 5, 33] it was shown that a modified version of the AKS sieve can approximate CVP in any essentially any norm to within a $(1 + \epsilon)$ factor in time and space $(1/\epsilon)^{O(n)}$ for any $\epsilon > 0$. Furthermore, these algorithms can solve CVP exactly in $2^{O(n)}$ time as long as the target point is “very close” to the lattice. It is worth noting that the AKS sieve is a *Monte Carlo* algorithm: while the output solution is correct with high probability, it is not guaranteed.

In a more recent breakthrough, Micciancio and Voulgaris [91] gave a *deterministic* $2^{O(n)}$ -time (and space) algorithm for exact SVP and CVP in the ℓ_2 norm, among many other lattice problems in NP. Interestingly, their algorithm works very differently from the AKS sieve, by computing an explicit description of the Voronoi cell of the lattice. (The Voronoi cell is the set of all points in \mathbb{R}^n that are closer to the origin than to any other lattice point.) In contrast to the AKS sieve, however, the algorithm of [91] appears to be quite specialized to ℓ_2 (or any norm defined by an ellipsoid, simply by applying a linear transformation). This is in part because in ℓ_2 the Voronoi cell is convex and has $2^{O(n)}$ facets, but in general norms this is not the case. A main problem left open in [91] was to find deterministic $2^{O(n)}$ -time algorithms for lattice problems in ℓ_p and other norms.

5.1.1 Lattice Problems

We define the lattice problems addressed in this chapter:

Definition 5.1.1 (Shortest Vector Problem). Given a $\mathbf{0}$ -centered convex body $K \subseteq \mathbb{R}^n$, and an n -dimensional lattice \mathcal{L} , the SVP with respect to K and \mathcal{L} is to compute an element of

$$\text{SVP}(K, \mathcal{L}) = \arg \min_{\mathbf{z} \in \mathcal{L} \setminus \{\mathbf{0}\}} \|\mathbf{z}\|_K.$$

Definition 5.1.2 (Closest Vector Problem). Given a $\mathbf{0}$ -centered convex body $K \subseteq \mathbb{R}^n$, an n -dimensional lattice \mathcal{L} , and target $\mathbf{x} \in \mathbb{R}^n$, the CVP with respect to K , \mathcal{L} and \mathbf{x} is to compute an element of

$$\text{CVP}(K, \mathcal{L}, \mathbf{x}) = \arg \min_{\mathbf{y} \in \mathcal{L}} \|\mathbf{y} - \mathbf{x}\|_K.$$

Given $\epsilon > 0$, the $(1 + \epsilon)$ -CVP with respect to K , \mathcal{L} and \mathbf{x} is to find $\mathbf{y} \in \mathcal{L}$ satisfying $\|\mathbf{y} - \mathbf{x}\|_K \leq (1 + \epsilon)d_K(\mathcal{L}, \mathbf{x})$ (distance from \mathbf{x} to \mathcal{L}).

We note that above definitions do not prescribe any computational assumptions on the set K . Since these details are technical and non-essential, we defer their presentation.

5.1.2 Results and Techniques

Our main contributions are new deterministic algorithms for solving SVP, CVP and $(1 + \epsilon)$ -CVP in general norms. These yield the first deterministic alternatives to the AKS randomized sieving approaches for these problems in norms other than ℓ_2 [2, 3, 18, 5, 33]. Compared to the AKS sieved based algorithms, our algorithms achieve similar running times (up to large $2^{O(n)}$ factors) and utilize less space (exactly $O(2^n)$). Furthermore, we show that if there exists an algorithm which solves CVP in ℓ_2 using $2^{O(n)}$ time and $S(n)$ space, then all our algorithms can be implemented using only $S(n)$ space. In particular, a polynomial space algorithm for ℓ_2 CVP would imply polynomial space versions of all our algorithms.

The first major tool underlying our algorithms is a new method for enumerating lattice points in any convex body. It uses as a crucial subroutine the Micciancio-Voulgaris (MV) algorithm [91] for the ℓ_2 norm that enumerates lattice points in an ellipsoid, and relies on the M-Ellipsoid concept from the previous chapter. This connection between lattice algorithms and convex geometry appears to be a fertile direction for further research.

For a lattice \mathcal{L} and convex body K in \mathbb{R}^n , let $G(K, \mathcal{L})$ be the largest number of lattice points contained in any translate of K , i.e.,

$$G(K, \mathcal{L}) = \max_{\mathbf{x} \in \mathbb{R}^n} |(K + \mathbf{x}) \cap \mathcal{L}|. \quad (5.1.1)$$

Our starting point is the following guarantee on the enumeration of $K \cap \mathcal{L}$.

Theorem 5.1.3 (Enumeration in convex bodies, informal). *Let $K \subseteq \mathbb{R}^n$ be a convex body and $\mathcal{L} \subseteq \mathbb{R}^n$ denote an n -dimensional lattice.*

- (1) *Given an ellipsoid $E \subseteq \mathbb{R}^n$, the set $K \cap \mathcal{L}$ can be outputted in deterministic time $G(K, \mathcal{L}) \cdot N(K, E) \cdot N(E, K) \cdot 2^{O(n)}$ using $O(2^n)$ space.*
- (2) *Using the deterministic M-Ellipsoid construction for K (Algorithm 4.1) to generate E , the set $K \cap \mathcal{L}$ can be outputted in time $G(K, \mathcal{L}) \cdot 2^{O(n)}$ using $O(2^n)$ space.*

From the above guarantees, we see that the running time of the lattice point enumeration procedure is minimized for an ellipsoid E where the product $N(K, E)N(E, K)$ is as small as possible. This shows us that a near optimal choice for E is in fact an M-Ellipsoid of K , which satisfies $N(E, K)N(K, E) = 2^{O(n)}$.

Definition 5.1.4 (Near Symmetric Norms). A general norm $\|\cdot\|_K$ is γ -symmetric, $0 < \gamma \leq 1$, if $\text{vol}(K) \geq \gamma^n \text{vol}(K \cap -K)$. $\|\cdot\|_K$ is “near-symmetric” if it is $\Omega(1)$ -symmetric.

For a centrally symmetric K , where $\|\cdot\|_K$ corresponds to a standard norm, K is 1-symmetric. If the centroid of K is at the origin, then $\|\cdot\|_K$ is $\frac{1}{2}$ -symmetric (see Theorem 2.3.7).

Our enumeration algorithm is at the core of the all the following applications. We begin with the Shortest Vector Problem in *any* near-symmetric norm $\|\cdot\|_K$.

Theorem 5.1.5 (SVP in any norm, informal). *There is a deterministic algorithm that, given any near symmetric norm $\|\cdot\|_K$ and n dimensional lattice \mathcal{L} , finds a shortest non-zero vector in \mathcal{L} under $\|\cdot\|_K$ in $2^{O(n)}$ time and $O(2^n)$ space.*

Here, the improvement over previous approaches is in the generalization to asymmetric norms defined by arbitrary convex bodies, as well as the deterministic nature of the algorithm.

We get a similar algorithm for the exact Closest Vector Problem, but its complexity grows with the distance from the target point to the lattice.

Theorem 5.1.6 (CVP in any norm, informal). *There is a deterministic algorithm that, given any near-symmetric norm $\|\cdot\|_K$, n dimensional lattice \mathcal{L} , and target $\mathbf{x} \in \mathbb{R}^n$, finds a closest vector to \mathbf{x} in \mathcal{L} under $\|\cdot\|_K$ in $(1 + 2\alpha)^n \cdot 2^{O(n)}$ time and $O(2^n)$ space, if the distance from \mathbf{t} to \mathcal{L} is at most $\alpha\lambda_1(K, \mathcal{L})$.*

To remove the dependence on the target distance, which in general is due to the existence of many short vectors which “confuse” the algorithm, we develop a method to “sparsify” any lattice while approximately maintaining its metric properties.

Theorem 5.1.7 (Lattice Sparsifier, informal). *There is a deterministic algorithm that, given any near-symmetric norm $\|\cdot\|_K$, n dimensional lattice \mathcal{L} , and distance $t \geq 0$, computes a sublattice $\mathcal{L}' \subseteq \mathcal{L}$ in deterministic $2^{O(n)}$ time and $O(2^n)$ space satisfying: (1) the distance from \mathcal{L}' to any point in the lattice is at most its distance to \mathcal{L} + t , (2) the number of points in \mathcal{L}' at distance t is at most $2^{O(n)}$.*

By combining the lattice sparsifier construction and our enumeration technique, we give a novel algorithm for solving the CVP approximately on any norm and lattice:

Theorem 5.1.8 (Approximate CVP in any norm, informal). *There is a deterministic algorithm that, given any near-symmetric norm in $\|\cdot\|_K$, n dimensional lattice \mathcal{L} , target $\mathbf{x} \in \mathbb{R}^n$, and $0 < \epsilon \leq 1$, finds a $(1 + \epsilon)$ -approximate closest vector to \mathbf{x} under $\|\cdot\|_K$ in \mathcal{L} , in $(1 + \frac{1}{\epsilon})^n \cdot 2^{O(n)}$ time and $O(2^n)$ space.*

For our next result, we show the utility of our general norm CVP solver, by using it to find a lattice point near the “center” of any given convex body. Though our algorithm is randomized, it is Las Vegas and uses only polynomial randomness.

Theorem 5.1.9 (Central Lattice Point, informal). *There is a randomized algorithm that, given any convex body K , n dimensional lattice \mathcal{L} , and $0 < \epsilon \leq 1$, returns $\mathbf{x} \in K$ such that $K - \mathbf{x}$ is $\frac{1}{5}$ -symmetric and a vector $\mathbf{y} \in \mathcal{L}$ such that $\|\mathbf{y} - \mathbf{x}\|_{K-\mathbf{x}} \leq (1 + \epsilon)d_{K-\mathbf{x}}(\mathcal{L}, \mathbf{x})$ in expected $(1 + \frac{1}{\epsilon})^n \cdot 2^{O(n)}$ time, using $O(2^n)$ space and $\text{poly}(n)$ randomness.*

One consequence of the above result is that it allows us to “approximately” decide whether K contains a point of \mathcal{L} , i.e. an approximate version of the Integer Programming problem (see Chapter 7). More precisely, the above algorithm allows us to distinguish between the following cases: either $K \cap \mathcal{L} \neq \emptyset$, or a $(1 + \epsilon)$ -scaling of K about \mathbf{x} contains no point of \mathcal{L} . To see this, if we run the algorithm with parameter ϵ , then if $K \cap \mathcal{L} \neq \emptyset$, we are guaranteed that $\|\mathbf{y} - \mathbf{x}\|_{K-\mathbf{x}} \leq (1 + \epsilon)$, and if $((1 + \epsilon)K - \epsilon\mathbf{x}) \cap \mathcal{L} = \emptyset$, then by definition $\|\mathbf{y} - \mathbf{x}\|_{K-\mathbf{x}} > (1 + \epsilon)$. Thus the two cases can be correctly distinguished as desired.

For our last result, we give a blackbox reduction from all of the above problems to ℓ_2 CVP. The following theorem states that the only thing required to reduce the space usage of our algorithms is a more space efficient algorithm for ℓ_2 CVP.

Theorem 5.1.10 (Reduction to ℓ_2 -CVP). *If there exists an algorithm for ℓ_2 CVP which runs in $2^{O(n)}$ time and $S(n)$ space (for $S(n)$ at least polynomial), then all the above algorithms can be implemented to run in the same time complexity while using only $S(n)$ space.*

In the rest of this introduction we give an overview of our enumeration technique and its application to SVP, CVP.

Enumeration via M-Ellipsoid coverings. We now explain the main technique underlying Theorem 5.1.3 (enumeration of lattice points in a convex body K). The idea is to reduce enumerating lattice points in a convex body K to enumerating lattice points inside translates of an M -ellipsoid E of K . First, for a translate $\mathbf{t} + E$, we show that a slight extension of the MV algorithm [91] for ℓ_2 CVP can be used to deterministically enumerate the points in $\mathbf{t} + E \cap \mathcal{L}$ in $2^{O(n)}G(E, \mathcal{L})$ time. From here, to enumerate the points in $K \cap \mathcal{L}$, we compute a covering Λ of K by E (i.e. $K \subseteq \Lambda + E$) and for each $\mathbf{t} \in \Lambda$ use the ellipsoid enumeration algorithm to compute $\mathbf{t} + E \cap \mathcal{L}$. This procedure computes a superset of $K \cap \mathcal{L}$, and so during the enumeration we simply ignore the lattice points that do not land in K . Finally, we do a small amount of extra processing to ensure that every lattice point in $K \cap \mathcal{L}$ is outputted exactly once.

For the complexity of the enumeration algorithm, we perform $|\Lambda| = 2^{O(n)}N(K, E)$ ellipsoid enumerations, each of which require $2^{O(n)}G(E, \mathcal{L})$ time. Since $G(E, \mathcal{L}) \leq N(E, K)G(K, \mathcal{L})$, the total running time is bounded by

$$2^{O(n)}N(K, E)G(E, \mathcal{L}) \leq 2^{O(n)}N(K, E)N(E, K)G(K, \mathcal{L}) \leq 2^{O(n)}G(K, \mathcal{L}). \quad (5.1.2)$$

as needed.

Shortest and Closest Vector Problems. Here we outline our deterministic $2^{O(n)}$ -time algorithm for SVP in any norm defined by a symmetric convex body K . (near-symmetric norms are dealt with similarly.)

Let \mathcal{L} be an n -dimensional lattice, and let $\lambda_1 = \lambda_1(K, \mathcal{L})$ be the length of its shortest vector under $\|\cdot\|_K$. We can assume by rescaling that $1/2 < \lambda_1 \leq 1$, so K contains an SVP solution. Our algorithm simply enumerates all nonzero points in $K \cap \mathcal{L}$ (using Theorem 5.1.3), and outputs one of the shortest. For the running time, it suffices to show that $G(K, \mathcal{L}) \leq 2^{O(n)}$, which follows by a simple packing argument: for any $x \in \mathbb{R}^n$, copies of $\frac{1}{4}K$ centered at each point in $(K + x) \cap \mathcal{L}$ are pairwise disjoint

(because $\lambda_1 > 1/2$) and contained in $\frac{5}{4}K + \mathbf{x}$, so $|(K + \mathbf{x}) \cap \mathcal{L}| \leq \text{vol}(\frac{5}{4}K)/\text{vol}(\frac{1}{4}K) = 5^n$.

For exact CVP with target point \mathbf{x} , the strategy is exactly the same as above, but we use a scaling dK so that $(dK - \mathbf{x}) \cap \mathcal{L} \neq \emptyset$ and $(\frac{d}{2}K - \mathbf{x}) \cap \mathcal{L} = \emptyset$ (i.e., d is a 2-approximation of the distance from \mathbf{x} to \mathcal{L}). In this case, the packing argument gives a bound of $G(dK, \mathcal{L}) \leq (1 + 2d/\lambda_1)^n$.

As we can see, the complexity of the exact CVP algorithm above depends on the distance of \mathbf{x} to the lattice. A natural question is whether such a dependence can be removed. In its exact version, the CVP (for near-symmetric norms) is at least as hard as IP, and so progress here seems hard. However, we show that if we are willing to tolerate near-optimal solutions, then our current enumeration framework can be modified to give $(1 + \epsilon)$ -approximate solution in single exponential time.

Given the above guarantees, the high level idea is straightforward: first preprocess \mathcal{L} to get a sublattice \mathcal{L}' , satisfying $d_K(\mathcal{L}', \mathbf{x}) \leq (1 + \epsilon)d_K(\mathcal{L})$, i.e. the distance from \mathcal{L}' to \mathbf{x} is approximately the same, and $G(d'K, \mathcal{L}') = 2^{O(n)}(1 + 1/\epsilon)^n$ (for $d' = d_K(\mathcal{L}', \mathbf{x})$), i.e. \mathcal{L}' is “sparse” at the target distance. Given such a \mathcal{L}' we simply run the current CVP on \mathcal{L}' and \mathbf{x} to get the desired solution. Hence the key for this approach is the construction of the “sparsifier” \mathcal{L}' for \mathcal{L} . For further details on the construction of lattice sparsifiers, we refer the reader to Section 5.5.

Interestingly, the above deterministic algorithms nearly match the running times (up to large $2^{O(n)}$ factors) of the randomized AKS sieve based algorithms for the same problems. Though our running times suffer from larger $2^{O(n)}$ factors compared to AKS due to our covering approach, we on the other hand save substantially on space. To contrast, the AKS sieving approach relies on sampling exponentially many “random” lattice points and combining them via a sieving procedure, until they all land in a desired convex region with overwhelming probability. For $(1 + \epsilon)$ -CVP, the AKS based approaches need to keep $2^{O(n)}(\frac{1}{\epsilon})^n$ in memory to perform the necessary sieving. In

our approach however, we use the Micciancio-Voulgaris techniques combined with an M-Ellipsoid cover to induce a graph on the lattice points we wish to enumerate, and then rely on a very space efficient graph traversal technique to output the desired lattice points. In particular, the space usage we require is proportional only to the maximum degree of the induced lattice point graph, which will be at most $2(2^n - 1)$ (corresponding the max number of voronoi relevant vectors of a lattice). Therefore, in contrast to AKS, our space usage does not depend on the number of outputted points.

5.2 *Lattice Point Enumeration in Convex Bodies*

In this section we prove our general enumeration theorem for convex bodies (Theorem 5.1.3, formalized in Theorem 5.2.6). This algorithm will be form the core of our algorithms for the Shortest and Closest Vector Problems (discussed in the following sections), as well as the Integer Programming Problem (discussed in Chapter 7).

To implement our general enumeration algorithm we will make crucial use of Micciancio and Voulgaris [91] algorithm for the Closest Vector Problem under the ℓ_2 norm (and general ellipsoidal norms), which we call the MV algorithm for short. In the following paragraphs, we discuss the key ideas behind the MV algorithm, as well as how it can be used to perform efficient lattice point enumeration inside an ellipsoid.

Voronoi Cell based Enumeration: For an ellipsoid $E = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_A \leq 1\}$, $A \succ 0$, define the voronoi cell of \mathcal{L} with respect to E as

$$\begin{aligned} \mathcal{V}(E, \mathcal{L}) &= \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_A \leq \|\mathbf{x} - \mathbf{y}\|_A, \forall \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\} \\ &= \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle_A \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle_A, \forall \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}. \end{aligned} \tag{5.2.1}$$

Here we remember that $\langle \mathbf{x}, \mathbf{y} \rangle_A = \mathbf{x}^t \mathbf{A} \mathbf{y}$ and that $\|\mathbf{x}\|_E = \|\mathbf{x}\|_A = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_A}$. We write \mathcal{V} for $\mathcal{V}(E, \mathcal{L})$ when the ellipsoid E and lattice \mathcal{L} are clear from context. From

its definition, we see that \mathcal{V} denotes the points in \mathbb{R}^n that are at least as close to $\mathbf{0}$ as any other lattice vector under $\|\cdot\|_E$. We note that \mathcal{V} is defined by infinitely many linear inequalities and is symmetric (since $\mathcal{L} = -\mathcal{L}$), and hence is a symmetric convex set. In the next theorem, we present the fundamental structural properties of \mathcal{V} .

Theorem 5.2.1 (Voronoi Cell Structure). *Let \mathcal{L} be an n dimensional, $E = E(A)$, $A \succ 0$, denote an ellipsoid in \mathbb{R}^n . Then $\mathcal{V} = \mathcal{V}(E, \mathcal{L})$ satisfies the following:*

- (1) $\mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x}) \Leftrightarrow \pm(\mathbf{y} - \mathbf{x}) \in \mathcal{V}$.
- (2) $\frac{1}{2}\lambda(E, \mathcal{L})E \subseteq \mathcal{V} \subseteq \mu(E, \mathcal{L})E$. Furthermore, \mathcal{V} is full dimensional and bounded.
- (3) $\text{int}(\mathcal{V})$ is \mathcal{L} -packing, and \mathcal{V} is \mathcal{L} -covering. Furthermore, $\text{vol}_n(\mathcal{V}) = \det(\mathcal{L})$.
- (4) \mathcal{V} is a full dimensional polytope.
- (5) The halfspace $\{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{x} \rangle_A \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle_A\}$, $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$, bounds a facet of \mathcal{V} iff $\text{CVP}(E, \mathcal{L}, \frac{1}{2}\mathbf{y}) = \{\mathbf{0}, \mathbf{y}\}$. Furthermore, \mathcal{V} has at most $2(2^n - 1)$ facets.

Proof.

Proof of 1. We note that $\mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x}) = \arg \min_{\mathbf{z} \in \mathcal{L}} \|\mathbf{z} - \mathbf{x}\|_E$ iff $\|\mathbf{y} - \mathbf{x}\| \leq \|\mathbf{z} - \mathbf{x}\|_E$ for all $\mathbf{z} \in \mathcal{L}$. Since $\mathbf{y} - (\mathcal{L} \setminus \{\mathbf{0}\}) = \mathbf{y} + (\mathcal{L} \setminus \{\mathbf{0}\}) = \mathcal{L} \setminus \{\mathbf{y}\}$, we can rewrite that last condition to $\|\mathbf{y} - \mathbf{x}\|_E \leq \|\mathbf{y} - \mathbf{v} - \mathbf{x}\|_E$ for all $\mathbf{v} \in \mathcal{L} \setminus \{\mathbf{0}\}$ (condition holds trivially for $\mathbf{v} = \mathbf{0}$), and hence by definition $\mathbf{v} - \mathbf{x} \in \mathcal{V}$. Lastly, since \mathcal{V} is symmetric we have that $\pm(\mathbf{v} - \mathbf{x}) \in \mathcal{V}$.

Proof of 2. Let $\lambda = \lambda_1(E, L)$. We will show that $\frac{1}{2}\lambda E \subseteq \mathcal{V}$. Assume not, there exists $\mathbf{x} \in \frac{1}{2}\lambda E$ such that $\mathbf{0}$ is not the closest lattice vector \mathbf{x} . Since $\|\mathbf{x} - \mathbf{0}\|_E = \|\mathbf{x}\|_E \leq \frac{1}{2}\lambda$, there must $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ satisfying $\|\mathbf{x} - \mathbf{y}\|_E < \frac{1}{2}\lambda$. But then note that

$$\|\mathbf{y}\|_E = \|\mathbf{y} - \mathbf{x} + \mathbf{x}\|_E \leq \|\mathbf{y} - \mathbf{x}\|_E + \|\mathbf{x}\|_E < \frac{1}{2}\lambda + \frac{1}{2}\lambda < \lambda$$

a clear contradiction. Let $\mu = \mu(E, \mathcal{L})$. We show that $\mathcal{V} \subseteq \mu E$. Assume not, then there exists $\mathbf{x} \in \mathcal{V}$ such that $\|\mathbf{x}\|_E > \mu$. But then, by definition of μ , there exists $\mathbf{y} \in \mathcal{L}$ such that $\|\mathbf{x} - \mathbf{y}\|_E \leq \mu$. But then \mathbf{y} is closer to \mathbf{x} than $\mathbf{0}$, a contradiction to \mathbf{x} being in \mathcal{V} .

Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ denote a basis for \mathcal{L} . From Lemma 2.4.7 we know that $\lambda \geq \min_i \|\mathbf{b}_i^*\|_{\pi_i(E)} > 0$ and that $\mu \leq \sum_{i=1}^n \|\mathbf{b}_i^*\|_{\pi_i(E)} < \infty$. Since $0 < \frac{1}{2}\lambda < \mu < \infty$, we have that \mathcal{V} is both full dimensional and bounded.

Proof of 3. We begin by showing that

$$\text{int}(\mathcal{V}) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_E < \|\mathbf{x} - \mathbf{y}\|_E, \forall \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}.$$

Clearly if $\mathbf{x} \in \text{int}(\mathcal{V})$, then none of the constraints of \mathcal{V} can be tight at \mathbf{x} , and hence $\|\mathbf{x}\|_E < \|\mathbf{x} - \mathbf{y}\|_E$ for all $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$. Now assume that $\|\mathbf{x}\|_E < \|\mathbf{x} - \mathbf{y}\|_E$ for all $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$. To show that $\mathbf{x} \in \text{int}(\mathcal{V})$, it suffices to show that for some $\epsilon > 0$, $\mathbf{x} + \epsilon E \subseteq \mathcal{V}$. Let $d = \inf\{\|\mathbf{y} - \mathbf{x}\|_E - \|\mathbf{x}\|_E : \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}$. We claim that $d > 0$. Take $\mathbf{y} \in \mathcal{L}$, and assume that $\|\mathbf{y}\|_E \geq 2\|\mathbf{x}\|_E + d + 1$. Then by the triangle inequality note that

$$\|\mathbf{y} - \mathbf{x}\|_E - \|\mathbf{x}\|_E \geq \|\mathbf{y}\|_E - 2\|\mathbf{x}\|_E \geq d + 1$$

From the above, we see that it suffices to evaluate the infimum over the points in $\mathcal{L} \setminus \{\mathbf{0}\} \cap (2\|\mathbf{x}\|_E + d + 1)E$. Since this set contains only a finite number of points, the infimum is achieved and hence $d > 0$ as needed. Letting $\epsilon = \frac{d}{2}$, we claim that $\mathbf{x} + \epsilon E \subseteq \mathcal{V}$. Take $\mathbf{z} \in \mathbf{x} + \epsilon E$. By the triangle inequality note that $\|\mathbf{z}\|_E \leq \|\mathbf{x}\|_E + \epsilon$. Furthermore, for $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ we have that

$$\|\mathbf{z} - \mathbf{y}\|_E = \|\mathbf{z} - \mathbf{x} + \mathbf{x} - \mathbf{y}\|_E \geq \|\mathbf{x} - \mathbf{y}\|_E - \|\mathbf{z} - \mathbf{x}\|_E \geq (\|\mathbf{x}\|_E + 2\epsilon) - \epsilon \geq \|\mathbf{z}\|_E.$$

Therefore $\mathbf{z} \in \mathcal{V}$ as needed.

We show that $\text{int}(\mathcal{V})$ is \mathcal{L} -packing. Now that for any $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$, $\|\mathbf{x} + \mathbf{y}\|_E > \|\mathbf{x}\|_E$ and hence $\mathbf{x} + \mathbf{y} \notin \mathcal{V}$. Therefore $|(\mathbf{x} + \mathcal{L}) \cap \text{int}(\mathcal{V})| = 1$, for all $\mathbf{x} \in \text{int}(\mathcal{V})$, as needed.

We now show that \mathcal{V} is \mathcal{L} -covering. Take $\mathbf{x} \in \mathbb{R}^n$, and let $\bar{\mathbf{x}} \in \arg \min_{\mathbf{z} \in \mathcal{L} + \mathbf{x}} \|\mathbf{z}\|_K$. Then by definition of $\bar{\mathbf{x}}$, we have that $\forall \mathbf{y} \in \mathcal{L}$ that $\|\bar{\mathbf{x}}\|_E \leq \|\bar{\mathbf{x}} + \mathbf{y}\|_E$. Therefore $\bar{\mathbf{x}} \in \mathcal{V}$, and hence $|(\mathcal{L} + \mathbf{x}) \cap \mathcal{V}| \geq 1$, as needed.

We prove that $\text{vol}_n(\mathcal{V}) = \det(\mathcal{L})$. By Lemma 2.4.4 since $\text{int}(\mathcal{V})$ is \mathcal{L} -packing we have $\text{vol}_n(\text{int}(\mathcal{V})) \leq \det(\mathcal{L})$, and since \mathcal{V} is \mathcal{L} -covering we have $\text{vol}_n(\mathcal{V}) \geq \det(\mathcal{L})$. The claim follows by noting that $\text{vol}_n(\text{int}(\mathcal{V})) = \text{vol}_n(\mathcal{V})$.

Proof of 4. We now show that \mathcal{V} is polyhedral. Let $S = 2\mu E \cap \mathcal{L}$. We claim that $\mathcal{V} = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle_A \ \forall \mathbf{y} \in S\}$. Since S is finite (as μE is bounded), this will prove polyhedrality of \mathcal{V} . To prove this, it suffices to show that if $\mathbf{x} \notin \mathcal{V}$, there exists $\mathbf{y} \in S$, such that $\|\mathbf{x} - \mathbf{y}\|_E < \|\mathbf{x}\|_E$. Now take $\mathbf{x} \notin \mathcal{V}$. Let $t = \min\{1, \frac{\mu}{\|\mathbf{x}\|_E}\}$ and let $\mathbf{w} = t\mathbf{x}$. Note that $0 \leq t \leq 1$ and that $\|\mathbf{w}\|_E \leq \min\{\mu, \|\mathbf{x}\|_E\} \leq \mu$. Assume that $\mathbf{w} \notin \mathcal{V}$. Then there exists $\mathbf{y} \in \mathcal{L}$, such that $\|\mathbf{w} - \mathbf{y}\|_E < \|\mathbf{w}\|_E \leq \mu$. Then note that

$$\|\mathbf{y}\|_E \leq \|\mathbf{y} - \mathbf{x}\|_E + \|\mathbf{x}\|_E \leq \mu + \mu = 2\mu$$

and hence $\mathbf{y} \in S$. Furthermore, remembering that $t\mathbf{x} = \mathbf{w}$, we have that

$$\begin{aligned} \|\mathbf{x} - \mathbf{y}\|_E &= \|(\mathbf{x} - \mathbf{w}) + (\mathbf{w} - \mathbf{y})\|_E \leq \|\mathbf{x} - \mathbf{w}\|_E + \|\mathbf{w} - \mathbf{y}\|_E \\ &< \|\mathbf{x} - \mathbf{w}\|_E + \|\mathbf{w}\|_E = \|(1-t)\mathbf{x}\|_E + \|t\mathbf{x}\|_E = \|\mathbf{x}\|_E \end{aligned}$$

as needed. Now assume that $\mathbf{w} \in \mathcal{V}$. Since $\mathbf{x} \notin \mathcal{V}$, we must have that $t < 1$ and hence $\|\mathbf{x}\|_E > \|\mathbf{w}\|_E = \mu$. Since $\mathcal{V} \subseteq \mu E$, we must have $\mathbf{w} \in \partial\mathcal{V}$, and hence there exist $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ such that $\|\mathbf{w} - \mathbf{y}\|_E = \|\mathbf{w}\|_E = \mu$. Using the exact same argument as above, we get that $\mathbf{y} \in S$ and that

$$\|\mathbf{x} - \mathbf{y}\|_E \leq \|\mathbf{x} - \mathbf{w}\|_E + \|\mathbf{w} - \mathbf{y}\|_E = \|(1-t)\mathbf{x}\|_E + \|t\mathbf{x}\|_E = \|\mathbf{x}\|_E$$

We show that the above inequality is strict. By Cauchy-Schwarz, we know that the above inequality is tight if and only if $\mathbf{x} - \mathbf{w} = s(\mathbf{w} - \mathbf{y})$ and $s > 0$. Hence $\mathbf{y} = (t + \frac{t-1}{s})\mathbf{x}$, and since $\|\mathbf{x} - \mathbf{y}\|_E = |(1-t)(\frac{s+1}{s})|\|\mathbf{x}\|_E = \|\mathbf{x}\|_E$, we must satisfy

$|(1-t)(\frac{s+1}{s})| = \pm 1$. Since $t < 1$ and $s > 0$, we must have that $s = \frac{1-t}{t} \Rightarrow \mathbf{y} = \mathbf{0}$, a contradiction to our assumption on \mathbf{y} . Hence the above inequality is strict, as needed.

Proof of 5. Take $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$, and let $t_y = \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle$. The halfspace $H_{\mathbf{y}, t_y}^{\leq} = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle_A \leq t_y\}$ (we use the inner product $\langle \cdot, \cdot \rangle_A$ for the rest of the proof) bounds a facet of \mathcal{V} if $H_{\mathbf{y}, t_y}^{\leq} \cap \mathcal{V}$ is $n-1$ dimensional.

Now assume that $\text{CVP}(E, \mathcal{L}, \frac{1}{2}\mathbf{y}) = \{\mathbf{0}, \mathbf{y}\}$. From here, we see that $\frac{1}{2}\mathbf{y} \in \mathcal{V}$, and that the only tight constraint at $\frac{1}{2}\mathbf{y}$ is $H_{\mathbf{y}, t_y}^{\leq}$. By the proof of (3), we now have that $\forall \mathbf{z} \in \mathcal{L} \setminus \{\mathbf{0}, \mathbf{y}\}$ that $\|\frac{1}{2}\mathbf{y} - \mathbf{z}\|_E \geq \|\mathbf{y}\|_E + \epsilon$ for some $\epsilon > 0$. Following the same argument as in (3), we get that $(\frac{1}{2}\mathbf{y} + \frac{\epsilon}{2}E) \cap H_{\mathbf{y}, t_y}^{\leq} \subseteq \mathcal{V}$, i.e. that $\mathcal{V} \cap H_{\mathbf{y}, t_y}^{\leq}$ is $n-1$ dimensional.

Now assume that $\text{CVP}(E, \mathcal{L}, \frac{1}{2}\mathbf{y}) \neq \{\mathbf{0}, \mathbf{y}\}$. Since the distance from $\|\frac{1}{2}\mathbf{y} - \mathbf{0}\|_E = \|\frac{1}{2}\mathbf{y} - \mathbf{y}\|_E$, there must exist $\mathbf{z} \in \text{CVP}(E, \mathcal{L}, \frac{1}{2}\mathbf{y}) \setminus \{\mathbf{0}, \mathbf{y}\}$. Therefore we have that

$$\|\frac{1}{2}\mathbf{y} - \mathbf{z}\|_E \leq \|\frac{1}{2}\mathbf{y}\|_E \Leftrightarrow \frac{1}{2} \langle \mathbf{y}, \mathbf{z} \rangle_A \geq \frac{1}{2} \langle \mathbf{z}, \mathbf{z} \rangle_A \Leftrightarrow \langle \mathbf{y} - \mathbf{z}, \mathbf{z} \rangle_A \geq 0$$

Now let $\mathbf{w} = \mathbf{y} - \mathbf{z}$, and let $t_z = \frac{1}{2} \langle \mathbf{z}, \mathbf{z} \rangle$ and $t_w = \frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle$. We claim that

$$H_{\mathbf{z}, t_z}^{\leq} \cap H_{\mathbf{w}, t_w}^{\leq} \subseteq H_{\mathbf{y}, t_y}^{\leq}$$

Assume that $\mathbf{x} \in H_{\mathbf{z}, t_z}^{\leq} \cap H_{\mathbf{w}, t_w}^{\leq}$. Since $\mathbf{w} + \mathbf{z} = \mathbf{y}$ and $\langle \mathbf{w}, \mathbf{z} \rangle_A \geq 0$, we have that

$$\begin{aligned} \langle \mathbf{y}, \mathbf{x} \rangle_A &= \langle \mathbf{w} + \mathbf{z}, \mathbf{x} \rangle_A \leq \frac{1}{2} (\langle \mathbf{w}, \mathbf{w} \rangle_A + \langle \mathbf{z}, \mathbf{z} \rangle_A) \leq \frac{1}{2} (\langle \mathbf{w}, \mathbf{w} \rangle_A + 2 \langle \mathbf{w}, \mathbf{z} \rangle_A + \langle \mathbf{z}, \mathbf{z} \rangle_A) \\ &= \frac{1}{2} \langle \mathbf{w} + \mathbf{z}, \mathbf{w} + \mathbf{z} \rangle_A = \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle_A \end{aligned}$$

as needed. First assume that $\langle \mathbf{z}, \mathbf{w} \rangle_A > 0$. From the above, this implies that $\mathbf{x} \in H_{\mathbf{z}, t_z}^{\leq} \cap H_{\mathbf{w}, t_w}^{\leq}$ satisfies $\langle \mathbf{x}, \mathbf{y} \rangle_A < \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle_A$. Since $\mathcal{V} \subseteq H_{\mathbf{z}, t_z}^{\leq} \cap H_{\mathbf{w}, t_w}^{\leq}$, this implies that $\mathcal{V} \cap H_{\mathbf{y}, t_y}^{\leq} = \emptyset$ and hence $H_{\mathbf{y}, t_y}^{\leq}$ does not bound a facet of \mathcal{V} . Now assume that $\langle \mathbf{z}, \mathbf{w} \rangle_A = 0$. Since $\mathbf{z} \notin \{\mathbf{0}, \mathbf{y}\}$, note that $\mathbf{w} \neq \mathbf{0}$, and hence $\langle \mathbf{z}, \mathbf{w} \rangle_A = 0 \Rightarrow \mathbf{z}, \mathbf{w}$ are linearly independent. Now take $\mathbf{x} \in \mathcal{V} \cap H_{\mathbf{y}, t_y}^{\leq}$. Since $\mathbf{x} \in H_{\mathbf{z}, t_z}^{\leq} \cap H_{\mathbf{w}, t_w}^{\leq}$,

$\mathbf{x} \in H_{\mathbf{y},t_y}^- \Rightarrow \mathbf{x} \in H_{\mathbf{z},t_z}^- \cap H_{\mathbf{w},t_w}^-$. Since \mathbf{y}, \mathbf{w} are linearly independent, we have that $\dim(\mathcal{V} \cap H_{\mathbf{y},t_y}^-) \leq \dim(H_{\mathbf{z},t_z}^- \cap H_{\mathbf{w},t_w}^-) = n - 2 \Rightarrow H_{\mathbf{y},t_y}^-$ does not bound a facet of \mathcal{V} .

We now show that \mathcal{V} has at most $2(2^n - 1)$ facets. First note that by symmetry of \mathcal{V} , $H_{\mathbf{y},t_y}^-$, $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$, bounds a facet of \mathcal{V} if and only if $H_{-\mathbf{y},t_y}^-$ bounds a facet of \mathcal{V} . Using the first part, we get the equivalence

$$\begin{aligned} H_{\pm\mathbf{y},t_y}^- \text{ bounds a facet of } \mathcal{V} &\Leftrightarrow \text{CVP}(E, \mathcal{L}, \frac{\mathbf{y}}{2}) = \{\mathbf{0}, \mathbf{y}\} \\ &\Leftrightarrow \arg \min_{\mathbf{z} \in \frac{\mathbf{y}}{2} + \mathcal{L}} \|\mathbf{z}\|_E = \left\{ \frac{\mathbf{y}}{2}, -\frac{\mathbf{y}}{2} \right\} \end{aligned}$$

Now note that $|\mathcal{L}/2 \pmod{\mathcal{L}}| = |\{\frac{1}{2}\mathbf{y} + \mathcal{L} : \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}| = 2^n - 1$ (see Lemma 2.4.2). Now since each pair $\{\mathbf{y}, -\mathbf{y}\} \in \mathcal{L} \setminus \{\mathbf{0}\}$ inducing facets of \mathcal{V} can be identified to a non-zero coset in $\mathcal{L}/2 \pmod{\mathcal{L}}$, we must have that number of facets is bounded by $2(2^n - 1)$ as needed. \square

The algorithm of Micciancio and Voulgaris make crucial use of the voronoi of the cell to solve CVP in the ℓ_2 norm (or any ellipsoidal norm). In particular, they show that the facet defining lattice vectors for $\mathcal{V}(E, \mathcal{L})$, which they denote the voronoi relevant vectors of \mathcal{L} with respect to E , form an extremely efficient ‘‘basis’’ for closest lattice vector search. Formally, we denote the set of voronoi relevant vectors as

$$\begin{aligned} \text{VR}(E, \mathcal{L}) = \{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\} : \text{the halfspace } \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle_A \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle_A\} \\ \text{is facet defining for } \mathcal{V}(E, \mathcal{L})\}. \end{aligned} \quad (5.2.2)$$

We write VR for $\text{VR}(E, \mathcal{L})$ whenever the context is clear. We now explain the utility of the voronoi relevant vectors for CVP. First, we note that since \mathcal{V} is a full dimensional polytope, \mathcal{V} is defined by its facet defining halfspaces, and hence $\mathcal{V} = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle_A \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle \mid \mathbf{y} \in \text{VR}\}$. Now assume we wish to solve CVP with respect \mathcal{L} and target $\mathbf{x} \in \mathbb{R}^n$ under $\|\cdot\|_E$. Let $\mathbf{y} \in \mathcal{L}$ denote the closest lattice vector to \mathbf{x} we have found thus far. From Theorem 5.2.1, we know that $\mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x}) \Leftrightarrow \mathbf{y} - \mathbf{x} \in \mathcal{V}$. Now if $\mathbf{y} \notin \text{CVP}(E, \mathcal{L}, \mathbf{x})$, given the description of \mathcal{V} , there exists $\mathbf{v} \in \text{VR}$

such that $\mathbf{y} - \mathbf{v} \in \mathcal{L}$ is closer to \mathbf{x} than \mathbf{y} . Hence if we have VR stored as a list, then a simple finite algorithm to compute a closest lattice vector to \mathbf{x} is as follows: start with any initial vector $\mathbf{y} \in \mathcal{L}$ and iterate moving to $\mathbf{y} + \mathbf{v}$, for $\mathbf{v} \in \text{VR}$ such that $\|\mathbf{y} - \mathbf{x}\|_E > \|\mathbf{y} + \mathbf{v} - \mathbf{x}\|_E$, as long as such an improving \mathbf{v} exists.

A refinement of the above procedure, which chooses the improving $\mathbf{v} \in \text{VR}$ more carefully, is essentially what is used in the MV algorithm to compute closest vectors. We note that the set of voronoi relevant vectors for \mathcal{L} must be computed in advance for it to be used in the above procedure. Indeed, much of the work performed by the MV algorithm is spent reducing the computation of the set VR to $2^{O(n)}$ CVPs on a lower dimensional sublattice \mathcal{L}' of \mathcal{L} , for which the set of voronoi relevant vectors has been precomputed. We refer the interested reader to [91] for the full details.

We now give a formal statement of the MV algorithm:

Theorem 5.2.2 (MV Algorithm [91]). *Let $E = E(A)$, $A \succ \mathbf{0}$, $A \in \mathbb{Q}^{n \times n}$, denote an ellipsoid in \mathbb{R}^n , \mathcal{L} denote an n -dimensional lattice with basis $B \in \mathbb{Q}^{n \times n}$. Then in 2^{2n} poly(\cdot) time and 2^n poly(\cdot) space, the MV algorithm can compute any of the following:*

- (1) *The set of voronoi relevant vectors $\text{VR}(E, \mathcal{L})$.*
- (2) *A shortest non-zero lattice vector $\mathbf{y} \in \text{SVP}(E, \mathcal{L})$.*
- (3) *Linearly independent vectors $\mathbf{y}_1, \dots, \mathbf{y}_n \in \mathcal{L}$ satisfying $\|\mathbf{y}_i\|_E = \lambda_i(E, \mathcal{L})$.*
- (4) *For a target vector $\mathbf{x} \in \mathbb{Q}^n$, a closest lattice vector $\mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x})$.*

Enumeration in an Ellipsoid: As a first task, we will show that the MV algorithm can be adapted to efficiently enumerate the lattices points inside an arbitrary ellipsoid. Given an ellipsoid $E = E(A)$, $A \succ 0$, shift \mathbf{x} and a lattice \mathcal{L} , we want to enumerate the set $(E + \mathbf{x}) \cap \mathcal{L}$.

To achieve this we first define the graph G with vertex set $V[G] = (E + \mathbf{x}) \cap \mathcal{L}$ and edge set $E[G] = \{\{\mathbf{y}, \mathbf{w}\} : \mathbf{y}, \mathbf{w} \in V, \mathbf{y} - \mathbf{w} \in \text{VR}(E, \mathcal{L})\}$. The idea will be to traverse the graph G using Avis-Fukuda reverse search ¹ (see section 4.2 for a thorough exposition) to enumerate the lattice points in $(E + \mathbf{x}) \cap \mathcal{L}$ in a time and space efficient manner. To traverse the graph using reverse search, we will need to build a good local search function f on V having a unique sink. Here we will show that a slight refinement of the “naive” CVP algorithm, which uses locally improving moves indexed by voronoi relevant vectors in $\text{VR}(E, \mathcal{L})$, can be used to give a local search function whose unique sink is a lexicographically minimal closest lattice vector to \mathbf{x} .

Let $<_{\text{lex}}$ denote the standard strict lexicographic ordering on \mathbb{R}^n , i.e. for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $\mathbf{x} <_{\text{lex}} \mathbf{y}$, if either $\mathbf{x}_1 < \mathbf{y}_1$ or $\mathbf{x}_1 = \mathbf{y}_1$ and $(\mathbf{x}_2, \dots, \mathbf{x}_n) <_{\text{lex}} (\mathbf{y}_2, \dots, \mathbf{y}_n)$.

We now present the local search function f on G and its associated guarantees:

Lemma 5.2.3. *Let $\mathbf{v}_1, \dots, \mathbf{v}_N$ denote an ordering of vectors in $\text{VR}(E, \mathcal{L})$. Let $f : \mathcal{L} \rightarrow \mathcal{L}$ denote the function where $f(\mathbf{y}) = \mathbf{y} - \mathbf{v}_i$, for a minimum index $i \in [N]$ such that either $\|\mathbf{y} - \mathbf{v}_i - \mathbf{x}\|_E < \|\mathbf{y} - \mathbf{x}\|_E$ or $\|\mathbf{y} - \mathbf{v}_i - \mathbf{x}\|_E = \|\mathbf{y} - \mathbf{x}\|_E$ and $\mathbf{y} - \mathbf{v}_i <_{\text{lex}} \mathbf{y}$, and $f(\mathbf{y}) = \mathbf{y}$ if no such i exists. Then we have that*

- (1) *f is a local search function on G , i.e. the graph $T(f) = \{(\mathbf{y}, f(\mathbf{y})) : \mathbf{y} \in V, f(\mathbf{y}) \neq \mathbf{y}\}$ is a directed acyclic subgraph of G .*
- (2) *The unique sink of f is the lexicographically minimal element of $\text{CVP}(E, \mathcal{L}, \mathbf{x})$.*
- (3) *f runs in $2^n \text{poly}(\cdot)$ time and space.*

Proof. We show that f is a valid local search function on G . By the definition of f , for a non-sink node $\mathbf{y} \in \mathcal{L}$ (i.e. $f(\mathbf{y}) \neq \mathbf{y}$) we have that either $\|f(\mathbf{y}) - \mathbf{x}\|_E < \|\mathbf{y} - \mathbf{x}\|_E$ or $\|f(\mathbf{y}) - \mathbf{x}\|_E = \|\mathbf{y} - \mathbf{x}\|_E$ and $f(\mathbf{y}) <_{\text{lex}} \mathbf{y}$. Therefore for a non-sink node $\mathbf{y} \in \mathcal{L}$, we have that $f^{(k)}(\mathbf{y}) \neq \mathbf{y}$ for all $k \geq 1$, and hence the graph $T(f)$ is directed acyclic.

¹We are indebted to Matthias Köppe for suggesting the use of reverse search to reduce the space complexity of our covering and enumeration algorithms.

Furthermore if $\mathbf{y} \in E + \mathbf{x}$ (i.e. $\|\mathbf{y} - \mathbf{x}\|_E \leq 1$), we clearly have that $f(\mathbf{y}) \in E + \mathbf{x}$. Since by construction of f , $\mathbf{y} - f(\mathbf{y}) \in \text{VR}(E, \mathcal{L})$, we get that $\{\mathbf{y}, f(\mathbf{y})\}$ is an edge of G . Therefore the graph $T(f)$ is a directed acyclic subgraph of G as needed.

We now show that $f(\mathbf{y}) = \mathbf{y}$, for $\mathbf{y} \in \mathcal{L}$, then \mathbf{y} is the lexicographically minimal element of $\text{CVP}(E, \mathcal{L}, \mathbf{x})$. Since $f(\mathbf{y}) = \mathbf{y}$, by construction of f we must have that $\|\mathbf{y} - \mathbf{x}\|_E \leq \|\mathbf{y} - \mathbf{v}_i - \mathbf{x}\|_E$ for all $i \in [N]$, and hence by Theorem 5.2.1 we have that $\mathbf{y} - \mathbf{x} \in \mathcal{V}(E, \mathcal{L}) \Leftrightarrow \mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x})$. Let \mathbf{y}' denote the lexicographically minimal element of $\text{CVP}(E, \mathcal{L}, \mathbf{x})$.

For the sake of contradiction, assume that $\mathbf{y}' \neq \mathbf{y}$. Since \mathcal{V} is symmetric, note $\mathbf{y} - \mathbf{x} \in \mathcal{V} \Rightarrow \mathbf{x} - \mathbf{y} \in \mathcal{V} \Rightarrow \mathbf{0} \in \text{CVP}(E, \mathcal{L}, \mathbf{x} - \mathbf{y})$. Since $\mathbf{y} \in \mathcal{L}$, $\text{CVP}(E, \mathcal{L}, \mathbf{x} - \mathbf{y}) = \text{CVP}(E, \mathcal{L}, \mathbf{x}) - \mathbf{y}$ and hence $\mathbf{y}' - \mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x} - \mathbf{y})$, and $\mathbf{y}' - \mathbf{y} \neq \mathbf{0}$. Letting $\mathbf{w} = \mathbf{y}' - \mathbf{y}$, noting that $\langle \cdot, \cdot \rangle_{\text{lex}}$ is translation invariant (i.e. $\mathbf{a} + \mathbf{c} \langle \cdot, \cdot \rangle_{\text{lex}} \mathbf{b} + \mathbf{c} \Leftrightarrow \mathbf{a} \langle \cdot, \cdot \rangle_{\text{lex}} \mathbf{b}$), we have that \mathbf{w} is the lexicographically minimal element of $\text{CVP}(E, \mathcal{L}, \mathbf{x} - \mathbf{y})$. Furthermore, since $\mathbf{0}, \mathbf{w} \in \text{CVP}(E, \mathcal{L}, \mathbf{x} - \mathbf{y})$, $\mathbf{w} \langle \cdot, \cdot \rangle_{\text{lex}} \mathbf{0}$ and

$$\|(\mathbf{x} - \mathbf{y}) - \mathbf{w}\|_E = \|\mathbf{x} - \mathbf{y}\|_E \Leftrightarrow \langle \mathbf{x} - \mathbf{y}, \mathbf{w} \rangle_A = \frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle_A.$$

Since $\mathbf{w} \in \mathcal{L} \setminus \{\mathbf{0}\}$, the inequality $\langle \cdot, \mathbf{w} \rangle_A \leq \frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle_A$ is a non-trivial valid inequality for \mathcal{V} , where we remember that

$$\mathcal{V} = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{v}_i \rangle_A \leq \frac{1}{2} \langle \mathbf{v}_i, \mathbf{v}_i \rangle \quad i \in [N]\} \quad (\text{since } \{\mathbf{v}_1, \dots, \mathbf{v}_N\} = \text{VR}(E, \mathcal{L})).$$

Therefore by Farkas Lemma, there exists non-negative multipliers $a_1, \dots, a_N \in \mathbb{R}_+$ such that

$$\sum_{i=1}^N a_i \mathbf{v}_i = \mathbf{w} \quad \text{and} \quad \sum_{i=1}^N \frac{1}{2} a_i \langle \mathbf{v}_i, \mathbf{v}_i \rangle_A \leq \frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle_A$$

Letting $I = \{i \in [N] : a_i > 0\}$, we see that

$$\frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle_A = \langle \mathbf{x} - \mathbf{y}, \mathbf{w} \rangle_A = \sum_{i \in I} a_i \langle \mathbf{x} - \mathbf{y}, \mathbf{v}_i \rangle \leq \sum_{i \in I} \frac{1}{2} a_i \langle \mathbf{v}_i, \mathbf{v}_i \rangle_A \leq \frac{1}{2} \langle \mathbf{w}, \mathbf{w} \rangle_A$$

Therefore all the inequalities above are equalities, and hence for all $i \in I$,

$$\langle \mathbf{x} - \mathbf{y}, \mathbf{v}_i \rangle = \frac{1}{2} \langle \mathbf{v}_i, \mathbf{v}_i \rangle_A \Rightarrow \|\mathbf{x} - \mathbf{y}\|_E = \|\mathbf{v}_i - (\mathbf{x} - \mathbf{y})\|_E.$$

Theorem 5.2.4 (Enumeration in Ellipsoids). *Algorithm 5.1 (Ellipsoid Enum) is correct and runs in $2^{O(n)} \cdot (1 + |(E + \mathbf{x}) \cap \mathcal{L}|) \cdot \text{poly}(\cdot)$ time using $2^n \text{poly}(\cdot)$ space.*

Proof.

Correctness: Using the MV algorithm (see Theorem 5.2.2), we first find a lexicographically minimal closest lattice vector $\mathbf{y} \in \mathcal{L}$ to \mathbf{x} under $\|\cdot\|_E$.

If $\|\mathbf{y} - \mathbf{x}\|_E > 1$, we return \emptyset . Note that if $\mathbf{y} \notin E + \mathbf{x}$, and since \mathbf{y} is the closest such vector we indeed get that $(E + \mathbf{x}) \cap \mathcal{L} = \emptyset$, as needed.

If $\mathbf{y} \in E + \mathbf{x}$, we begin by applying the local search function to \mathbf{y} (i.e. $\mathbf{y} \leftarrow f(\mathbf{y})$), until we find the lexicographically minimal closest lattice vector in $\text{CVP}(E, \mathcal{L}, \mathbf{x})$. By Lemma 5.2.3, the unique sink of f is exactly the desired vector, and hence this step terminates correctly. Lastly, we perform a reverse search starting from \mathbf{y} of the graph G on vertex set $(E + \mathbf{x}) \cap \mathcal{L}$ with respect to the local search function f . The correctness of this step follows directly from the correctness of the reverse search algorithm (see Algorithm 4.1) and the local search function f (see Lemma 5.2.3).

Running Time: The call to the MV algorithm for computing the voronoi relevant vectors requires $2^{2n} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space. Throughout the rest of the algorithm, the space needed is simply the space needed to store the list of voronoi relevant vectors $\text{VR}(E, \mathcal{L})$ (as well as a couple extra lattice vectors), for use in both the adjacency oracle Adj and local search function f . This requires $N \text{poly}(\cdot) = 2^n \text{poly}(\cdot)$ space.

In the next step, we apply the reverse search function f to \mathbf{y} until we reach the unique sink of f . We claim that the number of loop iterations is bounded by 2^n . By construction of f , we know that $\|f(\mathbf{y}) - \mathbf{x}\|_E \leq \|\mathbf{x} - \mathbf{x}\|_E$. Since $\mathbf{y} \in \text{CVP}(E, \mathcal{L}, \mathbf{x})$ at the beginning of the loop, \mathbf{y} remains a closest lattice vector to \mathbf{x} during every loop iteration. Since f never cycles, to bound the number of iterations it suffices

to bound $|\text{CVP}(E, \mathcal{L}, \mathbf{x})|$. We claim that $|\text{CVP}(E, \mathcal{L}, \mathbf{x})| \leq 2^n$. Assume for the sake of contradiction that $|\text{CVP}(E, \mathcal{L}, \mathbf{x})| \geq 2^n + 1$. Since the number of distinct cosets of $\mathcal{L} \pmod{2\mathcal{L}} = 2^n$, by the pigeon hole principle there exists distinct $\mathbf{y}_1, \mathbf{y}_2 \in \text{CVP}(E, \mathcal{L}, \mathbf{x})$ such that $\mathbf{y}_1 \equiv \mathbf{y}_2 \pmod{2\mathcal{L}} \Rightarrow \frac{1}{2}(\mathbf{y}_1 + \mathbf{y}_2) \in \mathcal{L}$. Now we have that

$$\left\| \frac{1}{2}(\mathbf{y}_1 + \mathbf{y}_2) - \mathbf{x} \right\|_E = \left\| \frac{1}{2}(\mathbf{y}_1 - \mathbf{x}) + \frac{1}{2}(\mathbf{y}_2 - \mathbf{x}) \right\|_E \leq \frac{1}{2} \|\mathbf{y}_1 - \mathbf{x}\|_K + \frac{1}{2} \|\mathbf{y}_2 - \mathbf{x}\|_E = \|\mathbf{y}_1 - \mathbf{x}\|_K \quad (5.2.3)$$

By the strict convexity of $\|\cdot\|_E$, we note that the above inequality can only hold with equality if $\mathbf{y}_1 - \mathbf{x} = t(\mathbf{y}_2 - \mathbf{x})$ for $t \geq 0$. Since $\|\mathbf{y}_1 - \mathbf{x}\|_E = \|\mathbf{y}_2 - \mathbf{x}\|_E$, if $\mathbf{y}_1 - \mathbf{x} = t(\mathbf{y}_2 - \mathbf{x})$ then $t = 1$. But then $\mathbf{y}_1 - \mathbf{x} = \mathbf{y}_2 - \mathbf{x}$, which is not possible since $\mathbf{y}_1 \neq \mathbf{y}_2$ by assumption. Therefore inequality (5.2.3) holds strictly, and hence $\frac{1}{2}(\mathbf{y}_1 + \mathbf{y}_2) \in \mathcal{L}$ is closer to \mathbf{x} than \mathbf{y}_1 and \mathbf{y}_2 , a clear contradiction. Now since the loop executes at most 2^n times, and each invocation of f takes $2^n \text{poly}(\cdot)$ time, the whole loop requires $2^{O(n)} \text{poly}(\cdot)$ time to execute.

Lastly, we run the Reverse-Search algorithm on the graph G starting from \mathbf{y} . By the guarantees of the Reverse-Search algorithm, outputting the lattices points in $(E + \mathbf{x}) \cap \mathcal{L}$ requires at most $N|(E + \mathbf{x}) \cap \mathcal{L}|$ calls to the adjacency oracle Adj and local search function f . Since both Adj and f requires at most $2^n \text{poly}(\cdot)$ space and time on each invocation, the full running time is $2^{O(n)}N(|(E + \mathbf{x}) \cap \mathcal{L}| + 1) = 2^{O(n)}(|(E + \mathbf{x}) \cap \mathcal{L}| + 1)$, and the space usage is $2^n \text{poly}(\cdot)$ as needed. \square

We now give a blackbox reduction from lattice point enumeration in an ellipsoid to CVP in the ℓ_2 norm. From the implementation of Ellipsoid-Enum, we see that to perform ellipsoid enumeration it suffices to be able to (1) compute a closest lattice vector to the center of the ellipsoid and (2) efficiently enumerate the voronoi relevant vectors of the lattice. Clearly, (1) is by definition achieved by any CVP solver so we may focus on (2). We see that the above implementation currently calls the MV algorithm to compute $\text{VR}(E, \mathcal{L})$, and then stores these vectors as a list. From here, we note that (a) all the accesses to the list VR are sequential (i.e. we iterate

over all the elements in the list), and (b) replacing VR by any superset would still suffice. Therefore to reduce the space usage of Lattice-Enum, it suffices to give a space efficient algorithm that sequentially outputs a superset of VR. Lastly, given the implementation, we must guarantee that the list is enumerated in the same order upon each invocation, but this is easily achieved.

We now give a simple space efficient and blackbox reduction from enumeration of the voronoi relevant vectors to CVP in the ℓ_2 norm (or more precisely, in ellipsoidal norms). This will enable us to satisfy requirement (2) above, and hence to reimplement Ellipsoid-Enum to use only as much space as the used CVP solver. We note that this reduction is already implicit in [91], and so we simply make it explicit here.

Algorithm 5.2 Voronoi-Enum($E, \mathcal{L}, \text{CVP-ALG}$)

Input: Ellipsoid $E = E(A)$, $A \succ 0$, n dimensional lattice \mathcal{L} with basis $B \in \mathbb{Q}^{n \times n}$, and solver CVP-ALG for ℓ_2 -CVP.

Output: Outputs set V satisfying $\text{VR}(E, \mathcal{L}) \subseteq S$.

- 1: **for all** $\mathbf{a} \in \{0, 1\}^n \setminus \{\mathbf{0}\}$ **do**
 - 2: $\mathbf{x} \leftarrow B\mathbf{a}$.
 - 3: $\mathbf{v} \leftarrow \text{CVP-ALG}(E, 2\mathcal{L}, \mathbf{x})$.
 - 4: **output** $\pm(\mathbf{v} - \mathbf{x})$.
-

Theorem 5.2.5. *Let CVP-ALG be an algorithm for ℓ_2 -CVP which runs in $2^{O(n)} \text{poly}(\cdot)$ time and $S(n) \text{poly}(\cdot)$ space on n dimensional lattices. Then given an ellipsoid $E = E(A)$, and an n dimensional lattice \mathcal{L} with basis $B \in \mathbb{Q}^{n \times n}$, Voronoi-Enum outputs a superset of $\text{VR}(E, \mathcal{L})$ using at most $2^{O(n)} \text{poly}(\cdot)$ time and $S(n) \text{poly}(\cdot)$ space.*

Analysis of Voronoi-Enum.

Correctness: To show that the algorithm is correct, we simply need to show that the algorithm outputs a superset of $\text{VR}(E, \mathcal{L})$. Take $\mathbf{y} \in \text{VR}(E, \mathcal{L})$. Then by Theorem 5.2.1, we have that

$$\left\{ \frac{1}{2}\mathbf{y}, -\frac{1}{2}\mathbf{y} \right\} = \arg \min_{\frac{1}{2}\mathbf{y} + \mathcal{L}} \|\mathbf{y}\|_E \Leftrightarrow \{\mathbf{y}, -\mathbf{y}\} = \arg \min_{\mathbf{y} + 2\mathcal{L}} \|\mathbf{y}\|_E$$

Now let $\bar{\mathbf{a}} = B^{-1}\mathbf{y} \in \mathbb{Z}^n$. Now let \mathbf{a} denote the unique element of $\bar{\mathbf{a}} + 2\mathbb{Z}^n \cap \{0, 1\}^n$.

We note that

$$\bar{\mathbf{a}} - \mathbf{a} \in 2\mathbb{Z}^n \Leftrightarrow B(\bar{\mathbf{a}} - \mathbf{a}) \in 2\mathcal{L} \Leftrightarrow \mathbf{y} \equiv B\mathbf{a} \pmod{2\mathcal{L}}$$

Assume we are in the loop iteration of Voronoi-Enum corresponding to \mathbf{a} above.

We claim that Voronoi-Enum outputs $\pm\mathbf{y}$. Since $\mathbf{x} = B\mathbf{a}$ is in the same coset of $2\mathcal{L}$ as \mathbf{y} , we have that $\text{CVP}(E, 2\mathcal{L}, \mathbf{x}) = \{\mathbf{y} + \mathbf{x}, \mathbf{x} - \mathbf{y}\}$. Therefore by correctness of CVP-ALG, we must have that $\mathbf{v} \leftarrow \text{CVP-ALG}(E, 2\mathcal{L}, \mathbf{x})$ satisfies $\mathbf{v} \in \{\mathbf{y} + \mathbf{x}, \mathbf{x} - \mathbf{y}\}$. Therefore the output $\pm(\mathbf{v} - \mathbf{x}) = \pm\mathbf{y}$. Hence Voronoi-Enum outputs a superset of $\text{VR}(E, \mathcal{L})$ as needed.

Runtime: The runtime of the algorithm is the time needed to solve $2^n - 1$ CVPs using algorithm CVP-ALG. Since CVP-ALG runs in $2^{O(n)} \text{poly}(\cdot)$ time at each invocation, the runtime is bounded by $(2^n - 1)2^{O(n)} \text{poly}(\cdot) = 2^{O(n)} \text{poly}(\cdot)$ time as needed. Furthermore, the space needed to run in the algorithm is simply the space needed to iterate over $\{0, 1\}^n$, which is $O(n)$, and the space needed to run CVP-ALG, which is $S(n) \text{poly}(\cdot)$. Therefore the total space usage of the algorithm is $S(n) \text{poly}(\cdot)$ as needed. \square

Proof of Theorem 5.1.10 (Reduction to ℓ_2 -CVP). The space requirements of all the algorithms in this chapter are dominated by the space needed to run algorithm Ellipsoid-Enum. Using algorithm 5.2 (Voronoi-Enum), Ellipsoid-Enum can be re-implemented to use only $S(n) \text{poly}(\cdot)$ space while maintaining the same time complexity. Therefore the space usage of all the algorithms in this chapter can be reduced to $S(n) \text{poly}(\cdot)$ without increasing their running times. \square

We can now state our enumeration theorem, which formalizes Theorem 5.1.3 from the introduction.

Algorithm 5.3 Algorithm Lattice-Enum(K, \mathcal{L}, ϵ)

Input: An (\mathbf{a}_0, r, R) -centered convex body $K \subseteq \mathbb{R}^n$ given by a weak membership oracle O_K , a basis $B \in \mathbb{Q}^{n \times n}$ for a lattice \mathcal{L} , and $0 < \epsilon \leq 1$.

Output: $S \subseteq \mathcal{L}$ satisfying (5.2.4).

- 1: Let $E(A) \leftarrow$ M-Ellipsoid(K) (Algorithm 4.1).
 - 2: Strengthen oracle O_K to an oracle O_K^* , where O_K^* satisfies $O_K^*(\mathbf{x}, \epsilon) = 1$ if $\mathbf{x} \in K$ and $O_K^*(\mathbf{x}, \epsilon) = 0$ if $\mathbf{x} \notin K^\epsilon$.
 - 3: Let $P \leftarrow \frac{1}{\sqrt{n}} A^{-\frac{1}{2}} [-1, 1]^n$ (maximum volume parallelepiped in $E(A)$).
 - 4: **for all** $\mathbf{x} \in$ Parallelepiped-Tiling($K, P, \frac{1}{n}$) **do**
 - 5: **for all** $\mathbf{y} \in$ Ellipsoid-Enum($E, \mathcal{L}, \mathbf{x}$) **do**
 - 6: **if** $O_K^*(\mathbf{y}, \epsilon) = 1$ and $\mathbf{y} \in P$ **then**
 - 7: **output** \mathbf{y} .
-

Theorem 5.2.6 (Enumeration in convex bodies). *Given a (\mathbf{a}_0, r, R) -centered convex body $K \subseteq \mathbb{R}^n$, and an n dimensional lattice \mathcal{L} , Algorithm 5.3 (Lattice-Enum) outputs a $S \subseteq \mathcal{L}$ (each lattice point is outputted exactly once) satisfying*

$$K \cap \mathcal{L} \subseteq S \subseteq (K + \epsilon B_2^n) \cap \mathcal{L} \quad (5.2.4)$$

in deterministic $G(K, \mathcal{L}) \cdot 2^{O(n)}$ poly(\cdot) time using 2^n poly(\cdot) space.

In the above algorithm, we are forced to go back and forth between ellipsoids and parallelepipeds at different stages. At the beginning, we start by computing the M-Ellipsoid E . Then for the purposes of efficient covering, we switch from E to the inscribed parallelepiped P . Next, we wish to compute the lattice points inside each shift of P intersected with K . We note that since P tiles K , each lattice point in $K \cap \mathcal{L}$ is contained exactly one shift of P , and hence is outputted exactly once. To compute $(P + \mathbf{x}) \cap \mathcal{L}$ for a shift \mathbf{x} , we again switch back to ellipsoids (since the MV technology is tailored for ellipsoids), and enumerate all the lattice points inside $(E + \mathbf{x}) \cap \mathcal{L}$ while outputting only those contained in $P + \mathbf{x}$.

We note that the parallelepiped covering used is identical to the one in Algorithm Ellipsoid-Cover. We will reference the analysis therein when appropriate.

Proof.

Correctness: The correctness is straightforward. We first compute an M-Ellipsoid E of K using Algorithm 4.1. From here, we use Parellelepiped-Tiling to compute over a tiling Λ of P by K (i.e. $K \subseteq \Lambda + P$), where $P = \frac{1}{n}A^{-\frac{1}{2}}[-1, 1]^n \subseteq E$ is a maximum volume (half-open) inscribed parallelepiped in E . For each element $\mathbf{x} \in \Lambda$, we use Ellipsoid-Enum to enumerate the lattice points in $E + \mathbf{x}$. Since $P \subseteq E$, we have that $K \subseteq \Lambda + E$, and hence the lattice points in $K \cap \mathcal{L}$ form a subset of the lattice points enumerated by the calls to Ellipsoid-Enum. Finally, the algorithms filters so that only lattice vectors $\mathbf{y} \in (E + \mathbf{x}) \cap \mathcal{L}$ satisfying $O_K^*(\mathbf{y}, \epsilon) = 1$ and $\mathbf{y} \in P + \mathbf{x}$ are outputted. The first condition guarantees that any lattice point in K is outputted and that no lattice point outside of $K + \epsilon B_2^n$ is outputted. The second guarantee ensures that we only output lattice points contained inside the current intersecting tile $\mathbf{x} + P$. Since every lattice point is contained in exactly one intersecting tile, this allows us to guarantee that every lattice point is outputted exactly once (note that a lattice point may indeed be contained in multiple shifts of the ellipsoid E). This completes the proof of correctness.

Runtime: By Theorem 4.1.3, Algorithm M-Ellipsoid computes an M-Ellipsoid E of K in $2^{O(n)}$ poly(\cdot) time using poly(\cdot) space. Next, we build a strengthening O_K^* of the weak membership oracle for O_K . Since K is (\mathbf{a}_0, r, R) -centered, O_K^* can be implemented in polynomial time from O_K (see Lemma 4.3.3 of [56]).

For the outer loop, using the analysis of Ellipsoid-Cover, we iterate a tiling Λ of K by P size at most $(3(\sqrt{\frac{\pi\epsilon}{2}})(1 + o(1)))^n N(K, E) = 2^{O(n)} N(K, E)$, which takes at most $2^{O(n)} N(K, E)$ poly(\cdot) time and poly(\cdot) space. For the inner loop, for each $\mathbf{x} \in \Lambda$, Ellipsoid-Enum iterates over the lattice points in $(E + \mathbf{x}) \cap \mathcal{L}$. By the guarantees on Ellipsoid-Enum, iterating over $(E + \mathbf{x}) \cap \mathcal{L}$ requires at most $2^{O(n)}(1 + |(E + \mathbf{x}) \cap \mathcal{L}|)$ poly(\cdot) $\leq 2^{O(n)} G(E, \mathcal{L})$ poly(\cdot) time and 2^n poly(\cdot) space. Next we note that $G(E, \mathcal{L}) \leq N(E, K)G(K, \mathcal{L})$, i.e. the maximum number of lattice points

in any translation of E is bounded by the maximum number of lattice points in any translate of K times the minimum number of translates of K needed to cover E . Hence the total time for enumeration is bounded by

$$2^{O(n)}N(K, E)G(E, L) \text{poly}(\cdot) = 2^{O(n)}N(K, E)N(E, K)G(K, L) \text{poly}(\cdot).$$

Since E is an M-Ellipsoid for K , we have that $N(K, E)N(E, K) = 2^{O(n)}$, and hence we get the desired runtime. Furthermore, we require at most $2^n \text{poly}(\cdot)$ space to perform the enumeration (this space is needed by Ellipsoid-Enum), and hence the total space usage is $2^n \text{poly}(\cdot)$ as needed. \square

5.3 Shortest Vector Problem

Our main goal will be to use the above enumeration algorithm to solve the Shortest Vector Problem. The following gives a useful bound on $G(K, L)$ for a general convex body. We recall that a γ -symmetric convex body K , satisfies the relation $\text{vol}_n(K \cap -K) \geq \gamma^n \text{vol}(K)$.

Lemma 5.3.1. *Let $K \subseteq \mathbb{R}^n$ denote a γ -symmetric convex body and let \mathcal{L} denote an n -dimensional lattice. Then for $d > 0$ we have that*

$$G(dK, \mathcal{L}) \leq \gamma^{-n} \left(1 + \frac{2d}{\lambda_1(K \cap -K, \mathcal{L})} \right)^n. \quad (5.3.1)$$

Since $K \cap -K \subseteq K$, we have that $\lambda_1(K \cap -K, \mathcal{L}) \geq \lambda_1(K, \mathcal{L})$. From this we see that the above bound is stronger than if we replaced $K \cap -K$ by K . Next, we note γ above is easily bounded in many natural situations. When K is centrally symmetric we can set $\gamma = 1$ since $K \cap -K = K$, and if K is a general convex body with $\mathbf{b}(K) = \mathbf{0}$ setting $\gamma = \frac{1}{2}$ is valid by Theorem 4.3.6.

Proof of Lemma 5.3.1. Let $s = \frac{1}{2}\lambda_1(K \cap -K, \mathcal{L})$. For $\mathbf{x} \in \mathcal{L}$, we examine

$$\mathbf{x} + \text{int}(s(K \cap -K)) = \{\mathbf{z} \in \mathbb{R}^n : \|\mathbf{z} - \mathbf{x}\|_{K \cap -K} < s\}.$$

Now for $\mathbf{x}, \mathbf{y} \in \mathcal{L}$, $\mathbf{x} \neq \mathbf{y}$, we claim that

$$\mathbf{x} + \text{int}(s(K \cap -K)) \cap \mathbf{y} + \text{int}(s(K \cap -K)) = \emptyset \quad (5.3.2)$$

Assume not, then $\exists \mathbf{z} \in \mathbb{R}^n$ such that $\|\mathbf{z} - \mathbf{x}\|_{K \cap -K}, \|\mathbf{z} - \mathbf{y}\|_{K \cap -K} < s$. Since $K \cap -K$ is symmetric, we note that $\|\mathbf{y} - \mathbf{z}\|_{K \cap -K} = \|\mathbf{z} - \mathbf{y}\|_{K \cap -K} < s$. But then we have that

$$\begin{aligned} \|\mathbf{y} - \mathbf{x}\|_{K \cap -K} &= \|\mathbf{y} - \mathbf{z} + \mathbf{z} - \mathbf{x}\|_{K \cap -K} \leq \|\mathbf{y} - \mathbf{z}\|_{K \cap -K} + \|\mathbf{z} - \mathbf{x}\|_{K \cap -K} \\ &< s + s = 2s = \lambda_1(K \cap -K, \mathcal{L}) \end{aligned}$$

a clear contradiction since $\mathbf{y} - \mathbf{x} \neq \mathbf{0}$.

Take $\mathbf{c} \in \mathbb{R}^n$. To bound $G(dK, \mathcal{L})$ we must bound $|(\mathbf{c} + dK) \cap \mathcal{L}|$. For $\mathbf{x} \in \mathbf{c} + dK$, we note that $\mathbf{x} + s(K \cap -K) \subseteq \mathbf{c} + (d + s)K$. Therefore,

$$\begin{aligned} \text{vol}_n((d + s)K) &= \text{vol}_n(\mathbf{c} + (d + s)K) \geq \text{vol}_n(((\mathbf{c} + dK) \cap \mathcal{L}) + s(K \cap -K)) \\ &= |(\mathbf{c} + dK) \cap \mathcal{L}| \text{vol}_n(s(K \cap -K)) \end{aligned}$$

where the last equality follows from (5.3.2). Therefore, we have that

$$|(\mathbf{c} + dK) \cap \mathcal{L}| \leq \frac{\text{vol}_n((d + s)K)}{\text{vol}_n(s(K \cap -K))} = \left(\frac{d + s}{\gamma s} \right)^n = \gamma^{-n} \left(1 + \frac{2d}{\lambda_1(K \cap -K, \mathcal{L})} \right)^n$$

as needed. □

We can now state the algorithm and main theorem of this section.

Theorem 5.3.2 (Analysis of Shortest-Vectors). *Let $\lambda_1 = \lambda_1(K, \mathcal{L})$. Given a γ -symmetric convex body $K \subseteq \mathbb{R}^n$, Algorithm 5.4 (Shortest-Vectors) outputs a set $S \subseteq \mathcal{L}$ satisfying*

$$\text{SVP}(K, \mathcal{L}) \subseteq S \subseteq \{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\} : \|\mathbf{y}\|_K \leq \lambda_1 + \epsilon \min\{1, \lambda_1\}\} \quad (5.3.3)$$

in deterministic $2^{O(n)}\gamma^{-n} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space.

Proof.

Algorithm 5.4 Shortest-Vectors(K, \mathcal{L}, ϵ)

Input: $(0, r, R)$ -centered convex body K with a weak distance oracle D_K for $\|\cdot\|_K$, a basis $B \in \mathbb{Q}^{n \times n}$ for a lattice \mathcal{L} , and $0 < \epsilon \leq 1$.

Output: Output $S \subseteq \mathcal{L}$ satisfying (5.3.3).

- 1: Compute $\mathbf{z} \in \text{SVP}(B_2^n, \mathcal{L})$ using the MV algorithm.
 - 2: $l \leftarrow \frac{\|\mathbf{z}\|_2}{R}; \epsilon_0 \leftarrow \frac{\epsilon}{3} \min\{1, l\}$
 - 3: $d \leftarrow \frac{l}{2}; \tilde{\lambda}_1 \leftarrow \infty$
 - 4: **repeat**
 - 5: $d \leftarrow 2d$
 - 6: **for all** $\mathbf{y} \in \text{Lattice-Enum}(dK, \mathcal{L}, r\epsilon_0)$ **do**
 - 7: **if** $\mathbf{y} \neq \mathbf{0}$ **then**
 - 8: $\tilde{\lambda}_1 \leftarrow \min\{\tilde{\lambda}_1, D_K(\mathbf{y}, \epsilon_0), d_f + \epsilon_0\}$
 - 9: **until** $\tilde{\lambda}_1 < \infty$
 - 10: **for all** $\mathbf{y} \in \text{Lattice-Enum}((\tilde{\lambda}_1 + \epsilon_0)dK, \mathcal{L}, r\epsilon_0)$ **do**
 - 11: **if** $\mathbf{y} \neq \mathbf{0}$ **then**
 - 12: **output** \mathbf{y}
-

Correctness: First note that since K is $(0, r, R)$ -centered, we know that $\frac{\|\mathbf{y}\|}{R} \leq \|\mathbf{y}\|_K \leq \frac{\|\mathbf{y}\|}{r}$ for all $\mathbf{y} \in \mathbb{R}^n$. Now take any $\mathbf{z} \in \text{SVP}(K, \mathcal{L})$ and $\tilde{\mathbf{z}} \in \text{SVP}(B_2^n, \mathcal{L})$. Here we note that $\lambda_1 = \|\mathbf{z}\|_K$. As in the algorithm, let $l = \frac{\|\tilde{\mathbf{z}}\|}{R}$. Now we see that

$$l = \frac{\|\tilde{\mathbf{z}}\|}{R} \leq \frac{\|\mathbf{z}\|}{R} \leq \|\mathbf{z}\|_K \leq \|\tilde{\mathbf{z}}\|_K \leq \frac{\|\tilde{\mathbf{z}}\|}{r} = l \frac{R}{r}$$

Therefore $l \leq \lambda_1 \leq l \frac{R}{r}$.

Let d_f denote the value of d after the first while loop terminates. We claim that $\frac{1}{2}d_f \leq \lambda_1 \leq d_f + \epsilon_0$. After the loop terminates, we are guaranteed that the call to $\text{Lattice-Enum}(d_f K, \mathcal{L}, r\epsilon_0)$ outputs a non-zero lattice vector $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$. Therefore, by the guarantees on Lattice-Enum

$$\mathbf{y} \in d_f K + r\epsilon_0 B_2^n \subseteq d_f K + \epsilon_0 K = (d_f + \epsilon_0)K$$

Since by definition $\|\mathbf{y}\|_K \geq \lambda_1$, we have that $\lambda_1 \leq d_f + \epsilon_0$. For the lower bound, we examine two cases. If the while loop finishes after the first iteration, we have that $\frac{1}{2}d_f < d_f = l \leq \lambda_1$ as needed. If the while loop executes more than once, we get that $\text{Lattice-Enum}(\frac{1}{2}d_f K, \mathcal{L}, r\epsilon_0)$ does not output any non-zero lattice points. By

the guarantees on Lattice-Enum, this implies that $\frac{1}{2}d_f K \cap \mathcal{L} = \{\mathbf{0}\} \Rightarrow \frac{1}{2}d_f \leq \lambda_1$ as needed.

We now claim that $\tilde{\lambda}_1$ (as in the algorithm) satisfies $\lambda_1 - \epsilon_0 \leq \tilde{\lambda}_1 \leq \lambda_1 + \epsilon_0$. We first note that $\tilde{\lambda}_1 = \min\{d_f + \epsilon_0, D_K(\mathbf{z}, \epsilon_0)\}$ from some non-zero $\mathbf{z} \in \mathcal{L}$. By the guarantees on D_K , we get that

$$\tilde{\lambda}_1 = \min\{d_f + \epsilon_0, D_K(\mathbf{z}, \epsilon_0)\} \geq \min\{\lambda_1, \|\mathbf{z}\|_K - \epsilon_0\} \geq \lambda_1 - \epsilon_0,$$

as needed. For the second inequality, we examine two cases. First assume that $\text{Lattice-Enum}(d_f K, \mathcal{L}, r\epsilon_0)$ outputs $\mathbf{z} \in \text{SVP}(K, \mathcal{L})$. Then $\tilde{\lambda}_1 \leq D_K(\mathbf{z}, \epsilon_0) \leq \lambda_1 + \epsilon_0$ as needed. If Lattice-Enum does not output any element of $\text{SVP}(K, \mathcal{L})$, we must have that $d_f < \lambda_1$ and hence $\tilde{\lambda}_1 \leq d_f + \epsilon_0 < \lambda_1 + \epsilon_0$, as needed.

We claim that if $\mathbf{y} \in \text{SVP}(K, \mathcal{L})$ then \mathbf{y} is outputted by the algorithm. Since $\|\mathbf{y}\|_K = \lambda_1$ and $\lambda_1 \leq \tilde{\lambda}_1 + \epsilon_0$, we have that $\mathbf{y} \in (\tilde{\lambda}_1 + \epsilon_0)K$. Therefore by the guarantees on Lattice-Enum, \mathbf{y} is outputted during the iteration over $\text{Lattice-Enum}((\tilde{\lambda}_1 + \epsilon_0)K, \mathcal{L}, r\epsilon_0)$ as needed. Lastly, any lattice point $\mathbf{y} \in \mathcal{L}$ outputted during the iteration satisfies $\mathbf{y} \neq \mathbf{0}$ and

$$\mathbf{y} \in (\tilde{\lambda}_1 + \epsilon_0)K + r\epsilon_0 B_2^n \subseteq (\tilde{\lambda}_1 + 2\epsilon_0)K \subseteq (\lambda_1 + 3\epsilon_0)K$$

Since $3\epsilon_0 = 3\frac{\epsilon}{3} \min\{1, l\} \leq \epsilon \min\{1, \lambda_1\}$, we have that $\|\mathbf{y}\|_K \leq \lambda_1 + \epsilon \min\{1, \lambda_1\}$ as needed.

Runtime: First, the computation of $\mathbf{y} \in \text{SVP}(B_2^n, \mathcal{L})$ using the MV algorithm takes 2^{2n} poly(\cdot) time and 2^n poly(\cdot) space. Next, we bound the number iterations of the while loop. If $d \geq \lambda_1$ during a loop iteration, by the guarantee on Lattice-Enum, the call to $\text{Lattice-Enum}(dK, \mathcal{L}, r\epsilon_0)$ is guaranteed to find a non-zero lattice vector. Since d starts at $l/2$, and by construction $\lambda_1 \leq \frac{R}{r}l$, we have that max number of loop iterations k , satisfies $2^k(l/2) \leq \frac{R}{r}l$, i.e. $k \leq \log_2(2\frac{R}{r})$.

By the proof of correctness, we have that $d_f \leq 2\lambda_1$, where d_f is the value of d at the last iteration. Therefore $d \leq 2\lambda_1$ at each loop iteration. Then by the

guarantees on Lattice-Enum, each call to $\text{Lattice-Enum}(dK, \mathcal{L}, r\epsilon_0)$ takes time at most $2^{O(n)}G(2\lambda_1 K, \mathcal{L}) \text{poly}(\cdot) = 2^{O(n)}\gamma^{-n} \text{poly}(\cdot)$ since K is γ -symmetric (Lemma 5.3.1). Since $\tilde{\lambda}_1 + \epsilon_0 \leq (1 + \epsilon)\lambda_1 \leq 2\lambda_1$, the complexity of the final call to Lattice-Enum is similarly bounded. Since the number of loop iterations $\log_2(2\frac{R}{r}) + 1$ is polynomial bounded, and each call to Lattice-Enum takes at most $2^{O(n)}\gamma^{-n} \text{poly}(\cdot)$ time, the total running time is $2^{O(n)}\gamma^{-n} \text{poly}(\cdot)$ as needed. Lastly, the space usage of all subroutines calls is at most $2^n \text{poly}(\cdot)$, and hence the total space usage is $2^n \text{poly}(\cdot)$ as needed. \square

5.4 Closest Vector Problem

Before presenting our CVP algorithm, we again need a simple enumeration bound.

Lemma 5.4.1. *Let $K \subseteq \mathbb{R}^n$ be a γ -symmetric convex body, and let $\mathcal{L} \subseteq \mathbb{R}^n$ denote an n -dimensional lattice. Then for $t > 0$ we have that*

$$G(tK, \mathcal{L}) \leq \gamma^{-n}(2t + 1)^n \cdot |(K \cap -K) \cap \mathcal{L}|$$

Furthermore, letting $\beta = \left(\max_{\mathbf{c} \in K} \frac{\text{vol}_n((K - \mathbf{c}) \cap (\mathbf{c} - K))}{\text{vol}_n(K)} \right)^{\frac{1}{n}}$, for $t > 0$ we have that

$$G(tK, \mathcal{L}) \leq \beta^{-n}(2t + 1)^n \cdot G(K, \mathcal{L}) \leq (4t + 2)^n \cdot G(K, \mathcal{L})$$

Proof. Examine $tK + \mathbf{x}$. Let $\mathbf{y}_1, \dots, \mathbf{y}_N \in (tK + \mathbf{x}) \cap \mathcal{L}$, denote a maximal collection of points such that the translates $\mathbf{y}_i + \frac{1}{2}(K \cap -K)$, $i \in [N]$, are interior disjoint. We claim that $(tK + \mathbf{x}) \cap \mathcal{L} \subseteq \cup_{i=1}^N \mathbf{y}_i + (K \cap -K)$. Take $\mathbf{z} \in (tK + \mathbf{x}) \cap \mathcal{L}$. Then by construction of $\mathbf{y}_1, \dots, \mathbf{y}_N$, there exists $i \in [N]$ such that

$$\mathbf{z} + \frac{1}{2}(K \cap -K) \cap \mathbf{y}_i + \frac{1}{2}(K \cap -K) \neq \emptyset \Rightarrow \mathbf{z} \in \mathbf{y}_i + (K \cap -K)$$

as needed. Therefore $|(tK + \mathbf{x}) \cap \mathcal{L}| \leq \sum_{i=1}^N |(\mathbf{y}_i + (K \cap -K)) \cap \mathcal{L}| = N|(K \cap -K) \cap \mathcal{L}|$.

Since K is γ -symmetric, we get that

$$N = \frac{\text{vol}_n(\cup_{i=1}^N \mathbf{y}_i + \frac{1}{2}(K \cap -K))}{\text{vol}_n(\frac{1}{2}(K \cap -K))} \leq 2^n \gamma^{-n} \frac{\text{vol}_n(tK + \frac{1}{2}(K \cap -K))}{\text{vol}_n(K)} \leq \gamma^{-n}(2t + 1)^n$$

as needed. Since the above bound holds for all $\mathbf{x} \in \mathbb{R}^n$, we get that

$$G(tK, \mathcal{L}) \leq \gamma^{-n}(2t+1)^n \cdot |(K \cap -K) \cap \mathcal{L}| \text{ as needed.}$$

Since $G(K, \mathcal{L})$ is invariant under translations of K , we may center K so that $\text{vol}(K \cap -K) = \beta^n \text{vol}(K)$. By Theorem 2.3.7, we have that

$$\text{vol}((K - \mathbf{b}(K)) \cap (\mathbf{b}(K) - K)) \geq 2^{-n} \text{vol}(K),$$

where $\mathbf{b}(K)$ is then centroid of K , and hence $\beta \geq \frac{1}{2}$. Using the first part of the lemma, we now get that

$$G(tK, \mathcal{L}) \leq \beta^{-n}(2t+1)^n \cdot |(K \cap -K) \cap \mathcal{L}| \leq \beta^{-n}(2t+1)^n \cdot G(K, \mathcal{L}) \leq (4t+2)^n \cdot G(K, \mathcal{L})$$

as needed. □

We can now state the algorithm and main theorem of this section.

Algorithm 5.5 Closest-Vectors($K, \mathcal{L}, \mathbf{x}, \epsilon$)

Input: $(0, r, R)$ -centered convex body K with weak distance oracle D_K for $\|\cdot\|_K$, a basis $B \in \mathbb{Q}^{n \times n}$ for a lattice \mathcal{L} , target point $\mathbf{x} \in \mathbb{Q}^n$, and $0 < \epsilon < 1$.

Output: $S \subseteq \mathcal{L}$ satisfying (5.4.1).

- 1: **if** $\mathbf{x} \in \mathcal{L}$ **then**
 - 2: **return** $\{\mathbf{x}\}$
 - 3: Compute $\mathbf{z} \in \text{CVP}(B_2^n, \mathcal{L}, \mathbf{x})$ using the MV algorithm.
 - 4: $l \leftarrow \frac{\|\mathbf{z} - \mathbf{x}\|}{R}$, $\epsilon_0 \leftarrow \frac{\epsilon}{3} \min\{1, l\}$
 - 5: $d \leftarrow \frac{l}{2}$, $\tilde{d}_x \leftarrow \infty$
 - 6: **repeat**
 - 7: $d \leftarrow 2d$
 - 8: **for all** $\mathbf{y} \in \text{Lattice-Enum}(dK + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ **do**
 - 9: $\tilde{d}_x \leftarrow \min\{\tilde{d}_x, D_K(\mathbf{y} - \mathbf{x}, \epsilon_0), d + \epsilon_0\}$
 - 10: **until** $\tilde{d}_x < \infty$
 - 11: **return** $\text{Lattice-Enum}((\tilde{d}_x + \epsilon_0)K + \mathbf{x}, \mathcal{L}, r\epsilon_0)$.
-

Theorem 5.4.2 (Correctness of Closest-Vectors). *Given a γ -symmetric convex body $K \subseteq \mathbb{R}^n$, Algorithm 5.5 outputs a set $S \subseteq \mathcal{L}$ such that*

$$\text{CVP}(K, L, \mathbf{x}) \subseteq S \subseteq \{\mathbf{y} \in \mathcal{L} : \|\mathbf{y} - \mathbf{x}\|_K \leq d_K(\mathcal{L}, \mathbf{x}) + \epsilon \min\{1, d_K(\mathcal{L}, \mathbf{x})\}\}. \quad (5.4.1)$$

in deterministic time $2^{O(n)}G(d_x K, \mathcal{L}) \text{poly}(\cdot)$ using at most $2^n \text{poly}(\cdot)$ space, where $d_x = d_K(\mathcal{L}, \mathbf{x})$. Furthermore if $d_x \leq \alpha \lambda_1(K \cap -K, \mathcal{L})$, the running time is bounded by $2^{O(n)} \gamma^{-n} (1 + 2\alpha)^n \text{poly}(\cdot)$.

The proof is essentially identical to the one for SVP.

Analysis of Closest-Vectors.

Correctness: First if $\mathbf{x} \in \mathcal{L}$, then \mathbf{x} is clearly the unique closest lattice vector to itself, and so we are done.

Next we note that since K is $(0, r, R)$ -centered, we know that $\frac{\|\mathbf{y}\|}{R} \leq \|\mathbf{y}\|_K \leq \frac{\|\mathbf{y}\|}{r}$ for all $\mathbf{y} \in \mathbb{R}^n$. Now take any $\mathbf{z} \in \text{CVP}(K, \mathcal{L}, \mathbf{x})$ and $\tilde{\mathbf{z}} \in \text{SVP}(B_2^n, \mathcal{L})$. Here we note that $d_x = \|\mathbf{z} - \mathbf{x}\|_K$. As in the algorithm, let $l = \frac{\|\tilde{\mathbf{z}} - \mathbf{x}\|}{R}$. Now we see that

$$l = \frac{\|\tilde{\mathbf{z}} - \mathbf{x}\|}{R} \leq \frac{\|\mathbf{z} - \mathbf{x}\|}{R} \leq \|\mathbf{z} - \mathbf{x}\|_K \leq \|\tilde{\mathbf{z}} - \mathbf{x}\|_K \leq \frac{\|\tilde{\mathbf{z}} - \mathbf{x}\|}{r} = l \frac{R}{r}$$

Therefore $l \leq d_x \leq l \frac{R}{r}$.

Let d_f denote the value of d after the first while loop terminates. We claim that $\frac{1}{2}d_f \leq d_x \leq d_f + \epsilon_0$. After the loop terminates, we are guaranteed that the call to $\text{Lattice-Enum}(d_f K + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ outputs a lattice vector $\mathbf{y} \in \mathcal{L}$. Therefore, by the guarantees on Lattice-Enum

$$\mathbf{y} \in d_f K + r\epsilon_0 B_2^n + \mathbf{x} \subseteq d_f K + \epsilon_0 K + \mathbf{x} = (d_f + \epsilon_0)K + \mathbf{x}$$

Since by definition $\|\mathbf{y} - \mathbf{x}\|_K \geq d_x$, we have that $d_x \leq d_f + \epsilon_0$. For the lower bound, we examine two cases. If the while loop finishes after the first iteration, we have that $\frac{1}{2}d_f < d_f = l \leq d_x$ as needed. If the while loop executes more than once, we get that $\text{Lattice-Enum}(\frac{1}{2}d_f K + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ does not output any lattice points. By the guarantees on Lattice-Enum , this implies that $(\frac{1}{2}d_f K + \mathbf{x}) \cap \mathcal{L} = \emptyset \Rightarrow \frac{1}{2}d_f \leq d_x$ as needed.

We now claim that \tilde{d}_x (as in the algorithm) satisfies $d_x - \epsilon_0 \leq \tilde{d}_x \leq d_x + \epsilon_0$. We first note that $\tilde{d}_x = \min\{d_f + \epsilon_0, D_K(\mathbf{z} - \mathbf{x}, \epsilon_0)\}$ from some $\mathbf{z} \in \mathcal{L}$. By the guarantees

on D_K , we get that

$$\tilde{d}_x = \min\{d_f + \epsilon_0, D_K(\mathbf{z} - \mathbf{x}, \epsilon_0)\} \geq \min\{d_x, \|\mathbf{z} - \mathbf{x}\|_K - \epsilon_0\} \geq d_x - \epsilon_0,$$

as needed. For the second inequality, we examine two cases. First assume that $\text{Lattice-Enum}(d_f K + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ outputs $\mathbf{z} \in \text{CVP}(K, \mathcal{L}, \mathbf{x})$. Then $\tilde{d}_x \leq D_K(\mathbf{z} - \mathbf{x}, \epsilon_0) \leq d_x + \epsilon_0$ as needed. If Lattice-Enum does not output any element of $\text{CVP}(K, \mathcal{L}, \mathbf{x})$, we must have that $d_f < d_x$ and hence $\tilde{d}_x \leq d_f + \epsilon_0 < d_x + \epsilon_0$, as needed.

We claim that if $\mathbf{y} \in \text{CVP}(K, \mathcal{L}, \mathbf{x})$ then \mathbf{y} is outputted by the algorithm. Since $\|\mathbf{y} - \mathbf{x}\|_K = d_x$ and $d_x \leq \tilde{d}_x + \epsilon_0$, we have that $\mathbf{y} \in (\tilde{d}_x + \epsilon_0)K + \mathbf{x}$. Therefore by the guarantees on Lattice-Enum , \mathbf{y} is outputted during the call to $\text{Lattice-Enum}((\tilde{d}_x + \epsilon_0)K + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ as needed. Lastly, any lattice point $\mathbf{y} \in \mathcal{L}$ outputted during the iteration satisfies

$$\mathbf{y} \in (\tilde{d}_x + \epsilon_0)K + r\epsilon_0 B_2^n + \mathbf{x} \subseteq (\tilde{d}_x + 2\epsilon_0)K + \mathbf{x} \subseteq (d_x + 3\epsilon_0)K + \mathbf{x}$$

Since $3\epsilon_0 = 3\frac{\epsilon}{3} \min\{1, l\} \leq \epsilon \min\{1, d_x\}$, we have that $\|\mathbf{y} - \mathbf{x}\|_K \leq d_x + \epsilon \min\{1, d_x\}$ as needed.

Runtime: First, we check whether $\mathbf{x} \in \mathcal{L}$, this takes polynomial time. Next, the computation of $\mathbf{z} \in \text{CVP}(B_2^n, \mathcal{L}, \mathbf{x})$ using the MV algorithm takes $2^{O(n)} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space. Next, we bound the number iterations of the while loop. If $d \geq d_x$ during a loop iteration, by the guarantee on Lattice-Enum , the call to $\text{Lattice-Enum}(dK + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ is guaranteed to find lattice vector. Since d starts at $l/2$, and by construction $d_x \leq \frac{R}{r}l$, we have that max number of loop iterations k , satisfies $2^k(l/2) \leq \frac{R}{r}l$, i.e. $k \leq \log_2(2\frac{R}{r})$.

By the proof of correctness, we have that $d_f \leq 2d_x$, where d_f is the value of d at the last iteration. Therefore $d \leq 2d_x$ at each loop iteration. Then by the guarantees on Lattice-Enum , each call to $\text{Lattice-Enum}(dK + \mathbf{x}, \mathcal{L}, r\epsilon_0)$ takes time at most $2^{O(n)}G(2d_x K, \mathcal{L}) \text{poly}(\cdot) = 2^{O(n)}G(d_x K, \mathcal{L}) \text{poly}(\cdot)$ by Lemma 5.4.1. Since

$\tilde{d}_x + \epsilon_0 \leq (1 + \epsilon)d_x \leq 2d_x$, the complexity of the final call to Lattice-Enum is similarly bounded. Since the number of loop iterations $\log_2(2\frac{R}{r}) + 1$ is polynomial bounded, and each call to Lattice-Enum takes at most $2^{O(n)}G(d_xK, \mathcal{L}) \text{poly}(\cdot)$ time, the total running time is $2^{O(n)}G(d_xK, \mathcal{L}) \text{poly}(\cdot)$ as needed. Lastly, the space usage of all subroutines calls is at most $2^n \text{poly}(\cdot)$, and hence the total space usage is $2^n \text{poly}(\cdot)$ as needed.

For the furthermore, if $d_x \leq \alpha\lambda_1(K \cap -K, \mathcal{L})$, then by γ -symmetry of K and Lemma 5.3.1 we have that $G(d_xK, \mathcal{L}) \leq \gamma^{-n}(1 + 2\alpha)^n$. The total running time is thus bounded by $2^{O(n)}\gamma^{-n}(1 + 2\alpha)^n \text{poly}(\cdot)$ as needed. \square

5.5 Approximate Closest Vector Problem

In this section, we give a single exponential algorithm to solve the approximate closest vector problem. In particular, we will give an algorithm which solves $(1 + \epsilon)$ -CVP under any norm in essentially $(1/\epsilon)^n$ time. To compare with Algorithm Closest-Vectors, we note that Closest-Vectors solves CVP under $\|\cdot\|_K$ exactly in $O(1/\epsilon)^n$ time as long as $d_K(\mathcal{L}, \mathbf{x}) \leq (1/\epsilon)\lambda_1(K, \mathcal{L})$. However, in the worst case $\max_{\mathbf{x} \in \mathbb{R}^n} d_K(\mathcal{L}, \mathbf{x}) = \mu(K, \mathcal{L}) \gg \lambda_1(K, \mathcal{L})$ (the ratio can be unbounded in fact), and hence the running time guarantee of Closest-Vectors is not in general meaningful.

To obtain our $(1 + \epsilon)$ -CVP algorithm, we shall essentially keep Closest-Vectors as the algorithmic base, however we shall first “sparsify” the input lattice before running it. Given a target vector \mathbf{x} , the high level idea is as follows. Instead of solving the CVP against \mathcal{L} , we solve it with respect to sublattice \mathcal{L}' which is (1) at approximately the same distance to \mathbf{x} as \mathcal{L} and (2) doesn’t have too many short vectors (i.e. is “sparse”). We formalize this as follows:

Definition 5.5.1 (Lattice Sparsifier). Let $K \subseteq \mathbb{R}^n$ be a γ -symmetric convex body, \mathcal{L} be an n -dimensional lattice, and $t \geq 0$. A sublattice $\mathcal{L}' \subseteq \mathcal{L}$ is a (K, t) sparsifier for \mathcal{L} if

$$(1) \quad \forall \mathbf{x} \in \mathbb{R}^n, d_K(\mathcal{L}', \mathbf{x}) \leq d_K(\mathcal{L}, \mathbf{x}) + t$$

$$(2) \quad G(tK, \mathcal{L}) = 2^{O(n)}\gamma^{-n}$$

To solve $(1 + \epsilon)$ -CVP on target \mathbf{x} , letting $d = d_K(\mathcal{L}, \mathbf{x})$, it suffices to solve the CVP on a $(K, \epsilon d)$ sparsifier \mathcal{L}' of \mathcal{L} . Here property (1) guarantees that \mathcal{L}' contains a $(1 + \epsilon)$ closest-vector, and (2) allows us to use lattice point enumeration to find this vector in roughly $(1 + 1/\epsilon)^n$ time. We note that \mathcal{L} is a (K, d) sparsifier of itself as long as $d = O(\lambda_1(K, \mathcal{L}))$.

We now give some simple equivalences for lattice sparsifiers:

Lemma 5.5.2. *Let K be a γ -symmetric convex body, and let \mathcal{L} be an n -dimensional lattice. Take $\mathcal{L}' \subseteq \mathcal{L}$ a full dimensional sublattice. For $t \geq 0$, we have that*

$$(1) \quad \forall \mathbf{y} \in \mathcal{L}, d_K(\mathcal{L}', \mathbf{y}) \leq t \Leftrightarrow \forall \mathbf{x} \in \mathbb{R}^n, d_K(\mathcal{L}', \mathbf{x}) \leq t + d_K(\mathcal{L}, \mathbf{x})$$

$$(2) \quad \mathcal{L}' \text{ is a } (K \cap -K, t) \text{ sparsifier} \Rightarrow \mathcal{L}' \text{ is a } (K, t) \text{ sparsifier.}$$

Proof. For statement (1), we prove the (\Rightarrow) implication (the other implication is immediate). Take $\mathbf{x} \in \mathbb{R}^n$, and pick $\mathbf{z} \in \text{CVP}(K, \mathcal{L}, \mathbf{x})$. Then by assumption, there exists $\mathbf{y} \in \mathcal{L}'$ such that $\|\mathbf{y} - \mathbf{z}\|_K \leq t$. Therefore

$$\begin{aligned} d_K(\mathcal{L}, \mathbf{x}) &\leq \|\mathbf{y} - \mathbf{x}\|_K = \|\mathbf{y} - \mathbf{z} + \mathbf{z} - \mathbf{x}\|_K \leq \|\mathbf{y} - \mathbf{z}\|_K + \|\mathbf{z} - \mathbf{x}\|_K \\ &\leq t + d_K(\mathcal{L}, \mathbf{x}), \end{aligned}$$

as needed.

We prove statement 2. Let $\mathcal{L}' \subseteq \mathcal{L}$ be a $(K \cap -K, t)$ sparsifier. Since $K \cap -K$ is 1-symmetric, by definition we have that $G(t(K \cap -K), \mathcal{L}') = 2^{O(n)}$. By Lemma 2.3.9 and γ -symmetry of K , we have that

$$N(tK, t(K \cap -K)) = N(K, K \cap -K) \leq 3^n \gamma^{-n}$$

Therefore

$$G(tK, \mathcal{L}') \leq G(t(K \cap -K), \mathcal{L}')N(tK, t(K \cap -K)) = 2^{O(n)}3^n \gamma^{-n} = 2^{O(n)}\gamma^{-n}.$$

Since $K \cap -K \subseteq K$, we have that $\|\mathbf{a}\|_K \leq \|\mathbf{a}\|_{K \cap -K}$ for all $\mathbf{a} \in \mathbb{R}^n$. Take $\mathbf{z} \in \mathcal{L}$. Since \mathcal{L}' is a $(K \cap -K, t)$ sparsifier there exists $\mathbf{y} \in \mathcal{L}'$ such that $t \geq \|\mathbf{y} - \mathbf{z}\|_{K \cap -K} \geq \|\mathbf{y} - \mathbf{z}\|_K$, as needed. Using part (1) of the lemma, we conclude that $\mathcal{L}' \subseteq \mathcal{L}$ is (K, t) sparsifier as needed. \square

From the above lemma, we see that it suffices to build lattice sparsifiers for symmetric convex bodies, i.e. to build a (K, t) sparsifier it suffices to build a $(K \cap -K, t)$ sparsifier for \mathcal{L} . Furthermore, the ‘‘closeness’’ requirement for a sparsifier of \mathcal{L} need only be checked against points in \mathcal{L} . Both these facts will help simplify the analysis.

The main technical contribution of this section is an explicit deterministic construction of a lattice sparsifier for any input norm and lattice. Apriori, it is unclear whether such sublattices should even exist even for simple norms such as ℓ_2 . As we will show, a simple probabilistic argument will guarantee the existence of such sublattices independent of the norm in question (see subsection 5.5.1).

The guarantees on our Lattice Sparsifier construction are as follows:

Theorem 5.5.3 (Algorithm Lattice-Sparsifier). *Let $K \subseteq \mathbb{R}^n$ be a $(\mathbf{0}, r, R)$ -centered and γ -symmetric convex body specified by a weak membership oracle O_K , and let \mathcal{L} denote an n dimensional lattice with a basis $B \in \mathbb{Q}^{n \times n}$. For $t \geq 0$, a (K, t) sparsifier for \mathcal{L} can be built using $2^{O(n)}$ $\text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space.*

Using the above construction, we first give a simple algorithm for $(1 + \epsilon)$ -CVP.

Theorem 5.5.4 (Correctness of Approx-Closest-Vectors). *For a γ -symmetric convex body $K \subseteq \mathbb{R}^n$, Algorithm 5.6 returns a non-empty set $S \subseteq \mathcal{L}$ of $(1 + \epsilon)$ -approximate closest vectors to \mathbf{x} in deterministic $2^{O(n)}\gamma^{-n}(1 + \frac{1}{\epsilon})^n \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space.*

The analysis is very similar to that of Closest-Vectors, so we reference the previous analysis whenever appropriate for the sake of concision.

Analysis of Approx-Closest-Vectors.

Algorithm 5.6 Approx-Closest-Vectors($K, \mathcal{L}, \mathbf{x}, \epsilon$)

Input: $(\mathbf{0}, r, R)$ -centered convex body $K \subseteq \mathbb{R}^n$ with weak distance oracle D_K for $\|\cdot\|_K$, a basis $B \in \mathbb{Q}^{n \times n}$ for \mathcal{L} , target $\mathbf{x} \in \mathbb{Q}^n$, $0 < \epsilon \leq 1$

Output: Outputs a non-empty set $S \subseteq \{\mathbf{y} \in \mathcal{L} : \|\mathbf{y} - \mathbf{x}\|_K \leq (1 + \epsilon)d_K(\mathcal{L}, \mathbf{x})\}$

```
1: if  $\mathbf{x} \in \mathcal{L}$  then
2:   return  $\{\mathbf{x}\}$ 
3: Compute  $\mathbf{z} \in \text{CVP}(B_2^n, \mathcal{L}, \mathbf{x})$  using the MV algorithm
4:  $l \leftarrow \frac{\|\mathbf{z} - \mathbf{x}\|_2}{R}$ ;  $\epsilon_0 \leftarrow \frac{\epsilon}{9} \min\{1, l\}$ 
5:  $d \leftarrow \frac{l}{2}$ ;  $\tilde{d}_x \leftarrow \infty$ 
6: repeat
7:    $d \leftarrow 2d$ 
8:    $\mathcal{L}' \leftarrow \text{Lattice-Sparsifier}(K, \mathcal{L}, \frac{\epsilon}{3}d)$ 
9:   for all  $\mathbf{y} \in \text{Lattice-Enum}((1 + \frac{\epsilon}{3})dK + \mathbf{x}, \mathcal{L}', r\epsilon_0)$  do
10:     $\tilde{d}_x \leftarrow \min\{\tilde{d}_x, D_K(\mathbf{y} - \mathbf{x}, \epsilon_0), (1 + \frac{\epsilon}{3})d + \epsilon_0\}$ 
11: until  $\tilde{d}_x < \infty$ 
12: return  $\text{Lattice-Enum}((\tilde{d}_x + \epsilon_0)K + \mathbf{x}, \mathcal{L}', r\epsilon_0)$ 
```

Correctness: If $\mathbf{x} \in \mathcal{L}$, we are clearly done. Next, by same analysis as Closest-Vectors, we get the guarantee that $l \leq d_x \leq l\frac{R}{r}$ where $d_x = d_K(\mathcal{L}, \mathbf{x})$.

Let d_f denote the value of d after the first while loop terminates. We claim that $\frac{1}{2}d_f \leq d_x \leq (1 + \epsilon/3)d_f + \epsilon_0$. When the while loop terminates, we are guaranteed that the call to $\text{Lattice-Enum}((1 + \frac{\epsilon}{3})d_f K + \mathbf{x}, \mathcal{L}', r\epsilon_0)$, outputs a lattice vector in \mathcal{L}' at distance at most $(1 + \frac{\epsilon}{3})d_f + \epsilon_0$ from \mathbf{x} . Since $\mathcal{L}' \subseteq \mathcal{L}$, we clearly have that $d_x \leq (1 + \frac{\epsilon}{3})d_f + \epsilon_0$ as needed.

If the while loop terminates after the first iteration, then $d_f = l \leq d_x$ and hence $\frac{1}{2}d_f < d_x$ as needed. If the loop iterates more than once, then for the sake of contradiction, assume that $\frac{1}{2}d_f > d_x$. Then in the before last iteration, the value of d is greater than d_x . Now we are guaranteed that $\text{Lattice-Sparsifier}(K, \mathcal{L}, \frac{\epsilon}{3}d)$ returns a lattice \mathcal{L}' satisfying

$$d_K(\mathcal{L}', \mathbf{x}) \leq d_K(\mathcal{L}, \mathbf{x}) + \frac{\epsilon}{3}d \leq (1 + \frac{\epsilon}{3})d$$

But then the call to $\text{Lattice-Enum}((1 + \frac{\epsilon}{3})dK + \mathbf{x}, \mathcal{L}', r\epsilon_0)$ is guaranteed to return a lattice point, and hence the while loop terminates at this iteration, a clear contradiction.

Hence $\frac{1}{2}d_f \leq d_x$ as needed.

Let $d'_x = d_K(\mathcal{L}', \mathbf{x})$, for \mathcal{L}' at the end of the while loop. By the same analysis as Closest Vectors, at the end loop we have that $d'_x - \epsilon_0 \leq \tilde{d}_x \leq d'_x + \epsilon_0$. Furthermore by construction of \mathcal{L}' , $d'_x \leq d_x + \epsilon/3d_f \leq (1 + 2\epsilon/3)d_x$.

Since $d'_x \leq \tilde{d}_x + \epsilon_0$, we know that $((\tilde{d}_x + \epsilon_0)K + \mathbf{x}) \cap \mathcal{L} \neq \emptyset$. Therefore we are guaranteed that the final call to $\text{Lattice-Enum}((\tilde{d}_x + \epsilon_0)K + \mathbf{x}, \mathcal{L}', r\epsilon_0)$ outputs all the closest vectors of \mathcal{L}' to \mathbf{x} . Lastly, any vector \mathbf{y} outputted during this call satisfies

$$\|\mathbf{y} - \mathbf{x}\|_K \leq \tilde{d}_x + 2\epsilon_0 \leq d'_x + 3\epsilon_0 \leq (1 + 2\epsilon/3)d_x + (\epsilon/3)l \leq (1 + \epsilon)d_x$$

as needed.

Running Time: We first bound the running time of each call to Lattice-Enum . Within the while loop, the calls to $\text{Lattice-Enum}((1 + \epsilon/3)dK + \mathbf{x}, \mathcal{L}', r\epsilon_0)$ run in $2^{O(n)}G((1 + \epsilon/3)dK, \mathcal{L}') \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space. By Lemma 5.4.1, since $(1 + \epsilon/3) = t(\epsilon/3)$ for $t = (3/\epsilon + 1)$, we have that

$$G((1 + \epsilon/3)dK, \mathcal{L}') \leq (4t + 2)^n G((\epsilon/3)d, \mathcal{L}') = 6^n (1 + 2/\epsilon)^n G((\epsilon/3)d, \mathcal{L}') = 2^{O(n)} \gamma^{-n} (1 + 1/\epsilon)^n$$

since by the guarantees on $\text{Lattice-Sparsifier}$, we have that $G((\epsilon/3)d, \mathcal{L}') = \gamma^{-n} 2^{O(n)}$.

Next the final call to $\text{Lattice-Enum}((\tilde{d}_x + \epsilon_0)K + \mathbf{x}, \mathcal{L}', r\epsilon_0)$ runs $2^{O(n)}G((\tilde{d}_x + \epsilon_0)K, \mathcal{L}') \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space. Now note that $\epsilon_0 \leq \frac{1}{9}\epsilon d_x$, and hence $(1 + \epsilon/3)d_f \geq d_x - \epsilon_0 \geq (1 - \epsilon/9)d_x$. From here we get that

$$d_f \geq \frac{1 - \epsilon/9}{1 + \epsilon/3} d_x \geq \frac{1 - 1/9}{1 + 1/3} d_x = 2/3 d_x$$

Finally, $\tilde{d}_x + \epsilon_0 \leq (1 + \epsilon/3)d_f + 2\epsilon_0 \leq (1 + \epsilon/3)d_f + 2/9\epsilon d_x \leq (1 + 2\epsilon/3)d_f$. Therefore, since $(1 + 2\epsilon/3) = t(\epsilon/3)$ for $t = (2 + 3/\epsilon)$, we get that

$$\begin{aligned} G((\tilde{d}_x + \epsilon_0)d_f K, \mathcal{L}') &\leq G((1 + 2\epsilon/3)d_f K, \mathcal{L}') \leq (4t + 2)^n G((\epsilon/3)d_f, \mathcal{L}') \\ &= (10 + 12/\epsilon)^n G((\epsilon/3)d_f, \mathcal{L}') = 2^{O(n)} \gamma^{-n} (1 + 1/\epsilon)^n \end{aligned}$$

by the guarantee on \mathcal{L}' .

Lastly, note that each call to Lattice-Sparsifier takes at most $2^{O(n)}$ poly(\cdot) time and 2^n poly(\cdot) space. Since the while loop iterates polynomially many times (i.e. at most $\log_2(2R/r)$), the total runtime is $2^{O(n)}\gamma^{-n}(1 + 1/\epsilon)^n$ poly(\cdot) and the total space usage is 2^n poly(\cdot) as needed. \square

The remainder of this section is dedicated to the lattice sparsifier construction.

5.5.1 A Simple Randomized Lattice Sparsifier Construction

We begin with an existence proof for lattice sparsifiers using the probabilistic method.

We will need the following classical sumset inequality:

Theorem 5.5.5. *Let $p \geq 1$ be a prime. Then for $A_1, \dots, A_k \subseteq \mathbb{Z}_p$, we have that*

$$|A_1 + \dots + A_k| \geq \min\left\{p, \sum_{i=1}^k |A_i| - k + 1\right\}$$

We will also need the following fact from number theory:

Theorem 5.5.6 (Bertrand's Postulate). *For every integer $k > 3$, there exists a prime $p \in \mathbb{Z}$ satisfying $k < p < 2k - 2$.*

We begin with the following crucial lemma. This forms the core of our lattice sparsifier construction.

Lemma 5.5.7. *Let $p \geq 5$ be a prime. Take $S \subseteq \mathbb{Z}_p^n$ satisfying $|S| < p < 2|S| - 2$ and $\mathbf{0} \in S$. Then there exists $\mathbf{a} \in \mathbb{Z}_p^n$ satisfying*

$$(1) \quad |\{\mathbf{y} \in S : \langle \mathbf{y}, \mathbf{a} \rangle \equiv 0 \pmod{p}\}| \leq 6$$

$$(2) \quad |\{\langle \mathbf{y}, \mathbf{a} \rangle \pmod{p} : \mathbf{y} \in S\}| \geq \frac{p}{4} + 1$$

Proof. Let \mathbf{a} denote a uniform random vector in \mathbb{Z}_p^n . The idea will be show that \mathbf{a} satisfies both conditions (1) and (2) with non-zero probability.

Let $N_0 = |S| - 1$. Let $E_i^{\mathbf{y}}$ denote the indicator for the event $\langle \mathbf{a}, \mathbf{y} \rangle \equiv i \pmod{p}$ for $\mathbf{y} \in S$ and $i \in \mathbb{Z}_p^n$.

Claim 1: $E[\sum_{\mathbf{y} \in S \setminus \{\mathbf{0}\}} E_0^{\mathbf{y}}] = N_0/p$

Proof. By linearity of expectation it suffices to prove that $E[E_0^{\mathbf{y}}] = \Pr[\langle \mathbf{a}, \mathbf{y} \rangle] = \frac{1}{p}$ for $\mathbf{y} \in S \setminus \{\mathbf{0}\}$. Since $\mathbf{y} \neq \mathbf{0}$, p is prime, and \mathbf{a} is uniform in \mathbb{Z}_p^n we have that $\langle \mathbf{a}, \mathbf{y} \rangle$ is uniform in \mathbb{Z}_p . Therefore $\Pr[\langle \mathbf{a}, \mathbf{y} \rangle] = \frac{1}{p}$ as needed. \square

Claim 2: Letting $E_i = \bigvee_{\mathbf{y} \in S} E_i^{\mathbf{y}}$, we have $E[\sum_{i \in \mathbb{Z}_p} E_i] \geq \frac{p-1}{p} N_0 (1 - \frac{N_0-1}{2p}) + 1$

Proof. Take $i \in \mathbb{Z}_p, i \neq 0$. For $\mathbf{y} \in S \setminus \{\mathbf{0}\}$, as $\langle \mathbf{a}, \mathbf{y} \rangle$ is uniform in \mathbb{Z}_p , we have that $\Pr[\langle \mathbf{w}, \mathbf{y} \rangle \equiv i \pmod{p}] = \frac{1}{p}$. Now take distinct $\mathbf{y}, \mathbf{z} \in S \setminus \{\mathbf{0}\}$, we claim that

$$\Pr[E_i^{\mathbf{y}} = 1 \wedge E_i^{\mathbf{z}} = 1] \leq \Pr[E_i^{\mathbf{y}} = 1] \Pr[E_i^{\mathbf{z}} = 1] = \frac{1}{p^2} \quad (5.5.1)$$

First assume that \mathbf{y}, \mathbf{z} are linearly dependent over \mathbb{Z}_p , i.e. $\mathbf{y} = a\mathbf{z}$, $a \in \mathbb{Z}_p$, $a \neq 1$, then if $\langle \mathbf{a}, \mathbf{y} \rangle \equiv i \pmod{p}$, we have that $\langle \mathbf{a}, \mathbf{y} \rangle \equiv a \langle \mathbf{a}, \mathbf{z} \rangle \equiv ai \pmod{p}$. Since $i \not\equiv 0 \pmod{p}$ and $a \not\equiv 1 \pmod{p}$, we must have that $ai \not\equiv i \pmod{p}$ and hence the probability of the event in (5.5.1) is 0. Next if \mathbf{y}, \mathbf{z} are linearly independent over \mathbb{Z}_p , then the random variables $\langle \mathbf{a}, \mathbf{y} \rangle$ and $\langle \mathbf{a}, \mathbf{z} \rangle$ are independent, and hence the inequality in (5.5.1) holds with equality. Now for $i \neq 0$, we have that

$$\begin{aligned} E[E_i] &= \Pr[\cup_{\mathbf{y} \in S \setminus \{\mathbf{0}\}} \{E_i^{\mathbf{y}} = 1\}] \geq \sum_{\mathbf{y} \in S \setminus \{\mathbf{0}\}} \Pr[E_i^{\mathbf{y}} = 1] - \sum_{\substack{\mathbf{y}, \mathbf{z} \in S \setminus \{\mathbf{0}\} \\ \mathbf{y} \neq \mathbf{z}}} \Pr[E_i^{\mathbf{y}} = 1 \wedge E_i^{\mathbf{z}} = 1] \\ &\geq \frac{1}{p} N_0 - \frac{1}{p^2} \binom{N_0}{2} = \frac{1}{p} N_0 (1 - \frac{N_0-1}{2p}) \end{aligned}$$

Next, since $\mathbf{0} \in S$, we have clearly have that $E_0 = 1$. Therefore, we get that

$$E[\sum_{i \in \mathbb{Z}_p} E_i] = E[\sum_{i \in \mathbb{Z}_p \setminus \{\mathbf{0}\}} E_i] + E[E_0] \geq \frac{p-1}{p} N_0 (1 - \frac{N_0}{2p}) + 1$$

as needed. \square

Let $\bar{E}_i = 1 - E_i$ for $i \in \mathbb{Z}_p$. Examine the events:

$$(1) B_1 : \sum_{\mathbf{y} \in S} E_0^{\mathbf{y}} \geq 7.$$

$$(2) B_2 : \sum_{i \in \mathbb{Z}_p} \bar{E}_i > \frac{3}{4}p - 1.$$

Notice that if the vector \mathbf{a} does not satisfy the conditions of the lemma, then either B_1 or B_2 must occur. Therefore it suffices to prove that $\Pr[B_1 \cup B_2] < 1$. By Markov's inequality, we have that

$$\Pr[B_1] = \Pr \left[\sum_{\mathbf{y} \in S} E_0^{\mathbf{y}} \geq 7 \right] = \Pr \left[\sum_{\mathbf{y} \in S \setminus \{\mathbf{0}\}} E_0^{\mathbf{y}} \geq 6 \right] \leq \frac{\mathbb{E} \left[\sum_{\mathbf{y} \in S \setminus \{\mathbf{0}\}} E_0^{\mathbf{y}} \right]}{6} \leq \frac{N_0}{6p} < \frac{1}{6}.$$

By our assumption on p , we know that

$$p < 2|S| - 2 = 2(N_0 + 1) - 2 \Rightarrow \frac{p}{2} < N_0 \Rightarrow \frac{p+1}{2} \leq N_0.$$

Since the lower bound on $\mathbb{E}[\sum_{i \in \mathbb{Z}_p} E_i]$ is an increasing function of N_0 (for $N_0 < p$), the bound is minimized for $N_0 = \frac{p+1}{2}$. From here, a straightforward computation reveals that

$$\begin{aligned} \mathbb{E} \left[\sum_{i \in \mathbb{Z}_p} E_i \right] &\geq \frac{p-1}{p} N_0 \left(1 - \frac{N_0-1}{2p} \right) + 1 \geq \frac{p-1}{p} \left(\frac{p+1}{2} \right) \left(1 - \frac{p-1}{4p} \right) + 1 \\ &= \frac{3}{8}p + \frac{9}{8} - \frac{3}{8p} - \frac{1}{8p^2} > \frac{3}{8}p + 1 \end{aligned}$$

for $p \geq 5$. From the above inequality we get that

$$\mathbb{E} \left[\sum_{i \in \mathbb{Z}_p} \bar{E}_i \right] < p - \frac{3}{8}p - 1 = \frac{5}{8}p - 1$$

Again by Markov's inequality, we have that

$$\Pr \left[\sum_{i \in \mathbb{Z}_p} \bar{E}_i > \frac{3}{4}p - 1 \right] \leq \frac{\mathbb{E} \left[\sum_{i \in \mathbb{Z}_p} \bar{E}_i \right]}{\frac{3}{4}p - 1} < \frac{\frac{5}{8}p - 1}{\frac{3}{4}p - 1} < \frac{\frac{5}{8}}{\frac{3}{4}} = \frac{5}{6}.$$

Hence $\Pr[B_1 \cup B_2] \leq \Pr[B_1] + \Pr[B_2] < \frac{1}{6} + \frac{5}{6} = 1$, as needed. \square

We now give our first lattice sparsifier construction.

Theorem 5.5.8. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body, $\mathcal{L} \subseteq \mathbb{R}^n$ be an n -dimensional lattice, and $t \geq 0$ be a non-negative number. Let $N = |tK \cap \mathcal{L}|$, and take p prime satisfying $N < p < 2N - 2$ if $N > 3$ and $p = 3$ otherwise. Then there exists $\mathbf{w} \in \mathcal{L}^*$ such that the sublattice $\mathcal{L}(\mathbf{w}) = \{\mathbf{y} \in \mathcal{L} : \langle \mathbf{w}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}$ satisfies*

$$(1) \forall \mathbf{x} \in \mathbb{R}^n, d_K(\mathcal{L}(\mathbf{w}), \mathbf{x}) \leq \mathcal{L} + 4t$$

$$(2) G(4tK, \mathcal{L}(\mathbf{w})) = 9^{n+1}$$

Proof. If $N \leq 3$, we let $\mathbf{w} = \mathbf{0}$, noting that $\mathcal{L}(\mathbf{0}) = \mathcal{L}$. Here condition (1) is trivially satisfied, and for condition (2) we have that

$$G(4tK, \mathcal{L}) \leq (2 \cdot 4 + 1)^n |tK \cap \mathcal{L}| \leq 9^n \cdot 3 \leq 9^{n+1}$$

by Lemma 5.4.1, as needed.

Now we assume that $N > 3$. By Bertrand's Postulate (i.e. theorem 5.5.6) there exists a prime p satisfying $N < p < 2N - 2$, as required by the theorem.

Claim 1: $p\mathcal{L} \cap 2tK = \{\mathbf{0}\}$.

Proof. For sake of contradiction, assume not and take $\mathbf{y} \in p\mathcal{L} \cap 2tK$, $\mathbf{y} \neq \mathbf{0}$. Then for $k \in \mathbb{Z}$, note that $(k/p)\mathbf{y} \in \mathcal{L}$ and

$$\|(k/p)\mathbf{y}\|_K = |k/p| \|\mathbf{y}\|_K \leq 2t |k/p|$$

by symmetry of K . Hence for $|k| \leq \lfloor p/2 \rfloor$, we get that $\|(k/p)\mathbf{y}\|_K \leq \frac{1}{2}2t = t$ and hence $(k/p)\mathbf{y} \in tK$. But then there at least $2\lfloor p/2 \rfloor + 1 \geq p > N$ distinct lattices points in $\mathcal{L} \cap tK$, a contradiction to our initial assumption. \square

Let $B^* = (\mathbf{b}^1, \dots, \mathbf{b}^n)$ denote a basis for \mathcal{L}^* . Let $S = \{B^{*T}\mathbf{y} \pmod{pZ^n} : \mathbf{y} \in tK \cap \mathcal{L}\}$. We claim that $|S| = |tK \cap \mathcal{L}| = N$. Assume not, then there exists distinct $\mathbf{y}_1, \mathbf{y}_2 \in tK \cap \mathcal{L}$ such that

$$B^{*T}\mathbf{y}_1 \equiv B^{*T}\mathbf{y}_2 \pmod{pZ^n} \Leftrightarrow B^{*T}(\mathbf{y}_1 - \mathbf{y}_2) \equiv \mathbf{0} \pmod{pZ^n} \Leftrightarrow \mathbf{y}_1 - \mathbf{y}_2 \in p\mathcal{L}$$

Now note that $\|\mathbf{y}_1 - \mathbf{y}_2\|_K \leq \|\mathbf{y}_1\|_K + \|\mathbf{y}_2\|_K \leq t + t = 2t$, and hence $\mathbf{y}_1 - \mathbf{y}_2 \in 2tK$. But then $\mathbf{y}_1 - \mathbf{y}_2 \in 2tK \cap \mathcal{L}$, in contradiction to Claim 1 since $\mathbf{y}_1 - \mathbf{y}_2 \neq \mathbf{0}$. Therefore $|S| = N$ as needed.

Since $\mathbf{0} \in S$ (since $\mathbf{0} \in tK$), and $|S| < p < 2|S| - 2$, by Lemma 5.5.7 there exists $\mathbf{a} \in \mathbb{Z}_p^n$ satisfying (a) $|\mathbf{y} \in S : \langle \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}| \leq 6$ and (b) $|\langle \mathbf{a}, \mathbf{y} \rangle \pmod{p} : \mathbf{y} \in S| \geq \frac{p}{4} + 1$. Let $\bar{\mathbf{a}}$ denote the unique representative of \mathbf{a} in $\{0, \dots, p-1\}^n$, and let $\mathbf{w} = B^* \bar{\mathbf{a}}$. The theorem will now follow directly from the following claim:

Claim 2: $\mathcal{L}(\mathbf{w})$ satisfies conditions (1) and (2).

Proof. Let $S_{in} = \{\mathbf{y} \in S : \langle \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}$ and let $C = \{\langle \mathbf{a}, \mathbf{y} \rangle \pmod{p} : \mathbf{y} \in S\}$. By our guarantees on \mathbf{a} , we know that $|S_{in}| \leq 6$ and $|C| \geq \frac{p}{4} + 1$.

We first prove establish condition (1). First, for $\mathbf{y} \in \mathcal{L}$ we have that

$$\mathbf{y} \in \mathcal{L}(\mathbf{w}) \Leftrightarrow \langle \mathbf{y}, \mathbf{w} \rangle \equiv 0 \pmod{p} \Leftrightarrow \langle \mathbf{y}, B^* \mathbf{a} \rangle \equiv 0 \pmod{p} \Leftrightarrow \langle B^{*T} \mathbf{y}, \mathbf{a} \rangle \equiv 0 \pmod{p}$$

$$\begin{aligned} \text{Therefore } \mathbf{y} \in tK \cap \mathcal{L}(\mathbf{w}) &\Leftrightarrow \mathbf{y} \in tK \cap \mathcal{L} \text{ and } \langle B^{*T} \mathbf{y}, \mathbf{a} \rangle \equiv 0 \pmod{p} \\ &\Leftrightarrow \mathbf{y} \in tK \cap \mathcal{L} \text{ and } B^{*T} \mathbf{y} \pmod{p\mathbb{Z}^n} \in S_{in}. \end{aligned}$$

Since the map $\mathbf{y} \rightarrow B^{*T} \mathbf{y} \pmod{p\mathbb{Z}^n}$ is injective restricted to $tK \cap \mathcal{L}$, we have that $|tK \cap \mathcal{L}(\mathbf{w})| = |S_{in}| \leq 6$. Now by Lemma 5.4.1, we have that

$$G(4tK, \mathcal{L}(\mathbf{w})) \leq 9^n \cdot |tK \cap \mathcal{L}(\mathbf{w})| \leq 9^n \cdot 6 \leq 9^{n+1},$$

as needed.

We now establish condition (2), i.e. for any $\mathbf{x} \in \mathbb{R}^n$, $d_K(\mathcal{L}(\mathbf{w}), \mathbf{x}) \leq d_K(\mathcal{L}, \mathbf{x}) + 4t$. By Lemma 5.5.2 it suffices to prove (2) for $\mathbf{x} \in \mathcal{L}$. Take $\mathbf{z} \in \mathcal{L}$. Let $c_{\mathbf{z}} \equiv \langle B^{*T} \mathbf{z}, \mathbf{a} \rangle \pmod{p}$. By Theorem 5.5.5, since $C \subseteq \mathbb{Z}_p$, $|C| \geq \frac{p}{4} + 1$, we have that

$$|C + C + C + C| \geq \min\{p, 4(\frac{p}{4} + 1) - 3\} \geq p$$

and hence $C + C + C + C = \mathbb{Z}_p$. Therefore there exists $\mathbf{y}_1, \dots, \mathbf{y}_4 \in tK \cap \mathcal{L}$ satisfying

$$c_{\mathbf{z}} \equiv \sum_{i=1}^4 \langle B^{*T} \mathbf{y}_i, \mathbf{a} \rangle \equiv \left\langle B^{*T} \sum_{i=1}^4 \mathbf{y}_i, \mathbf{a} \right\rangle \pmod{p}$$

Let $\mathbf{y} = \sum_{i=1}^4 \mathbf{y}_i \in \mathcal{L}$. Now we have that

$$\langle B^{*T} \mathbf{z}, \mathbf{a} \rangle \equiv c_{\mathbf{z}} \equiv \langle B^{*T} \mathbf{y}, \mathbf{a} \rangle \pmod{p} \Leftrightarrow \mathbf{z} - \mathbf{y} \in \mathcal{L}(\mathbf{w}),$$

and that $\|\mathbf{y}\|_K = \|\sum_{i=1}^4 \mathbf{y}_i\|_K \leq \sum_{i=1}^4 \|\mathbf{y}_i\|_K \leq 4t$. Therefore letting $\mathbf{v} = \mathbf{z} - \mathbf{y} \in \mathcal{L}(\mathbf{w})$, we get that

$$d_K(\mathcal{L}, \mathbf{z}) \leq \|\mathbf{v} - \mathbf{z}\|_K = \|\mathbf{y}\|_K \leq 4t$$

as needed. □

□

5.5.2 Derandomizing the Lattice Sparsifier Construction

We first summarize the lattice sparsifier construction from the previous section. Let K be a symmetric convex body, and \mathcal{L} be an n dimensional lattice. To build a (K, t) sparsifier, we do as follows

- (1) Compute $N \leftarrow |tK \cap \mathcal{L}|$, and prime p satisfying $N < p < 2N - 2$.
- (2) Build basis $B^* \in \mathbb{Q}^{n \times n}$ for \mathcal{L}^* and compute $S \leftarrow \{B^{*T} \mathbf{y} \pmod{p} : \mathbf{y} \in tK \cap \mathcal{L}\}$.
- (3) Find vector $\mathbf{a} \in \mathbb{Z}_p^n$ satisfying

$$(a) \quad |\{\mathbf{y} \in S : \langle \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}| \leq 6 \quad (b) \quad |\{\langle \mathbf{a}, \mathbf{y} \rangle : \mathbf{y} \in S\}| \geq \frac{p}{4} + 1$$

- (4) Return sublattice $\mathcal{L}' = \{\mathbf{y} \in \mathcal{L} : \langle \mathbf{y}, B^* \mathbf{a} \rangle \equiv 0 \pmod{p}\}$.

To implement the above construction efficiently and deterministically, we must overcome several obstacles. First the number of lattices points N in $tK \cap \mathcal{L}$ could be very large (since we have no control on t), and hence we cannot hope to compute N or the set S efficiently via lattice point enumeration. Second, the construction of the vector \mathbf{a} is probabilistic (see Lemma 5.5.7), and hence we will need to replace it with an explicit deterministic construction.

To overcome the first difficulty, we will build the (K, t) sparsifier iteratively. In particular, we will compute a sequence of sparsifiers $\mathcal{L}'_1, \dots, \mathcal{L}'_k$, satisfying that \mathcal{L}'_{i+1} is a $(K, c^i \lambda)$ sparsifier for \mathcal{L}'_i for $i \geq 0$, where $\mathcal{L}'_0 = \mathcal{L}$, $\lambda = \lambda_1(K, \mathcal{L})$ and $c > 1$ is a constant. Since we start the sparsification process at the minimum distance of \mathcal{L} , and only increase the sparsification distance by a constant factor at each step, we will be able to guarantee that the number of lattice points we process at each step is bounded by $2^{O(n)}$. Furthermore, the geometric growth rate in the sparsification distance will allow us to conclude that \mathcal{L}'_i is in fact a $(K, \frac{c^{i+1}}{c-1} \lambda)$ sparsifier for \mathcal{L} . Hence, iterating the process roughly $k \approx \ln \frac{t}{\lambda_1}$ steps will yield the final desired sparsifier.

For the second difficulty, i.e. the deterministic construction of \mathbf{a} , the main idea is to use a dimension reduction procedure² which allows \mathbf{a} to be computed efficiently via exhaustive enumeration (i.e. trying all possible \mathbf{a} 's). Let N and S be as in the description. Since $N < p < 2N - 2$, we note that exhaustive search over \mathbb{Z}_p^n requires a search over $p^n \leq (2N)^n$ possibilities, and the validity check (i.e. conditions (a) and (b)) for any particular \mathbf{a} can be implemented in $\text{poly}(N)$ time by simple counting. Since the existence of the desired \mathbf{a} depends only on $|S|$ (and not on n), if we can compute a linear projection $\pi : \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^{n-1}$ such that $\pi(S) = |S|$, then we can restrict to finding a good $\mathbf{a} \in \mathbb{Z}_p^{n-1}$ for $\pi(S)$. Indeed, we will show that such a map π can be computed efficiently and deterministically as long as $n \geq 3$. Therefore, repeating the process $n - 2$ times, we are left with finding a good $\mathbf{a} \in \mathbb{Z}_p^2$, which can do by trying all $p + 1 \leq 2N$ lines in \mathbb{Z}_p^2 . As discussed in the previous paragraph, we will be able to guarantee that $N = 2^{O(n)}$, and hence the entire construction described above will be implementable in $2^{O(n)}$ time and space as desired.

²We are indebted to Gabor Kuhn for suggesting the dimension reduction procedure in Algorithm 5.7

Algorithm 5.7 Algorithm Good-Vector(S, p)

Input: $S \subseteq \mathbb{Z}_p^n$, $\mathbf{0} \in S$, integer $n \geq 1$, p a prime satisfying $|S| < p < 2|S| - 2$.

Output: $\mathbf{a} \in \mathbb{Z}_p^n$ satisfying conditions of Lemma 5.5.7 .

- 1: **if** $n = 1$, **return** 1
 - 2: $P \leftarrow I_n$ ($n \times n$ identity)
 - 3: **for** n_0 **in** n **to** 3 **do**
 - 4: **for all** $\mathbf{q} \in \text{Lines}(\mathbb{Z}_p^{n_0})$ **do**
 - 5: Compute basis $B \in \mathbb{Z}_p^{n_0 \times n_0 - 1}$ satisfying $\mathbf{q}^\perp = B\mathbb{Z}_p^{n_0 - 1}$
 - 6: \forall distinct $\mathbf{x}, \mathbf{y} \in PS$ check that $B^T \mathbf{x} \not\equiv B^T \mathbf{y} \pmod{p\mathbb{Z}_p^{n_0 - 1}}$.
 If no collisions, set $P \leftarrow B^T P$ and exit loop; otherwise, continue.
 - 7: **for all** $\mathbf{a} \in \text{Lines}(\mathbb{Z}_p^2)$ **do**
 - 8: Compute $zeros \leftarrow |\{\mathbf{y} \in PS : \langle \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}|$
 - 9: Compute $distinct \leftarrow |\{\langle \mathbf{a}, \mathbf{y} \rangle \pmod{p} : \mathbf{y} \in PS\}|$
 - 10: **if** $zeros \leq 6$ and $distinct \geq \frac{p}{4} + 1$ **then**
 - 11: **return** $P^t \mathbf{a}$
-

We begin with the deterministic algorithm implementing Lemma 5.5.7. We denote the set of lines in \mathbb{Z}_p^n as

$$\text{Lines}(\mathbb{Z}_p^n) = \bigcup_{i=1}^n \{\mathbf{a} \in \mathbb{Z}_p^n : \mathbf{a}_j = 0, 1 \leq j < i, \mathbf{a}_i = 1, \mathbf{a}_k \in \mathbb{Z}_p, i < k \leq n\}.$$

It is easy to see that $\text{Lines}(\mathbb{Z}_p^n)$ contains a unique representative of every 1 dimensional subspace of \mathbb{Z}_p^n , and that $|\text{Lines}(\mathbb{Z}_p^n)| = \sum_{i=0}^{n-1} p^i = \frac{p^n - 1}{p - 1}$. For a vector $\mathbf{q} \in \mathbb{Z}_p^n$ we denote the orthogonal complement $\mathbf{q}^\perp = \{\mathbf{y} \in \mathbb{Z}_p^n : \langle \mathbf{q}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}$.

In the desired application of the above algorithm, the set S above will in fact be represented implicitly. Here the main access methodology we will require from S is a way to iterate over its elements. In the context of $(1 + \epsilon)$ -CVP, the enumeration method over S will correspond to the Lattice-Enum algorithm. Here we state the guarantees of the algorithm abstractly in terms of the number of iterations that required over S .

Theorem 5.5.9 (Good-Vector). *Algorithm 5.7 is correct, and performs $\text{poly}(n, \log p)p^4$ arithmetic operations as well as $O(np^3)$ iterations over the elements of S . Furthermore, the space usage (not counting the space needed to iterate over S) is $\text{poly}(n, \log p)$.*

Analysis of Good-Vector.

Correctness: We must show that the outputted vector \mathbf{a} satisfies the guarantees of Lemma 5.5.7, i.e. that for the returned vector $\mathbf{a} \in \mathbb{Z}_p^n$ satisfies:

$$(1) |\{\mathbf{y} \in S : \langle \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}| \leq 6$$

$$(2) |\{\langle \mathbf{a}, \mathbf{y} \rangle \pmod{p} : \mathbf{y} \in S\}| \geq \frac{p}{4} + 1$$

If $n = 1$, setting $\mathbf{a} \in \mathbb{Z}_p$ to 1 (i.e. line 1), trivially satisfies (1) and (2). We assume $n \geq 2$. We prove the following invariant for the first loop (line 2): at the beginning of each iteration, $P \in \mathbb{Z}_p^{n_0 \times n}$ and $|PS| = |S|$.

First let's assume that during the loop iteration, we find $B \in \mathbb{Z}_p^{n_0 \times (n_0-1)}$ satisfying the $B^T \mathbf{x} \neq B^T \mathbf{y}$ for all distinct $\mathbf{x}, \mathbf{y} \in PS$ (verified in line 5). From this condition we have that the map $\mathbf{x} \rightarrow B^T \mathbf{x}$ is injective when restricted to PS , and hence $|B^T PS| = |S|$. Next, since $B \in \mathbb{Z}_p^{n_0 \times (n_0-1)}$ and $P \in \mathbb{Z}_p^{n_0 \times n}$, we have that P is set to $B^T P \in \mathbb{Z}_p^{n_0-1 \times n}$ for the next iteration as needed.

Now we show that a valid projection matrix B^T is guaranteed to exist as long as $n_0 \geq 3$. First, we claim that there exists $\mathbf{q} \in \text{Lines}(\mathbb{Z}_p^{n_0})$, such that for all distinct $\mathbf{x}, \mathbf{y} \in PS$, $(\mathbb{Z}_p \mathbf{q} + \mathbf{x}) \cap (\mathbb{Z}_p \mathbf{q} + \mathbf{y}) = \emptyset$, i.e. all the lines passing through PS in the direction \mathbf{q} are disjoint. Since $\mathbf{y} - \mathbf{x} \neq \mathbf{0}$, \mathbb{Z}_p is a field, and $\mathbb{Z}_p \mathbf{q}$ is a subgroup of \mathbb{Z}_p^n , we have that

$$\begin{aligned} (\mathbb{Z}_p \mathbf{q} + \mathbf{x}) \cap (\mathbb{Z}_p \mathbf{q} + \mathbf{y}) \neq \emptyset &\Leftrightarrow \mathbb{Z}_p \mathbf{q} + (\mathbf{x} - \mathbf{y}) \cap \mathbb{Z}_p \neq \emptyset \\ &\Leftrightarrow \mathbf{x} - \mathbf{y} \in \mathbb{Z}_p \mathbf{q} \Leftrightarrow \mathbf{q} \in \mathbb{Z}_p(\mathbf{x} - \mathbf{y}). \end{aligned}$$

Therefore a line $\mathbb{Z}_p \mathbf{q}$ fails to satisfy (a) if and only if \mathbf{q} is contained in the line $\mathbb{Z}_p(\mathbf{x} - \mathbf{y})$ for distinct $\mathbf{x}, \mathbf{y} \in PS$. Clearly, the number of lines that can be generated in this way from PS is at most $\binom{|PS|}{2} = \binom{|S|}{2} < \frac{p(p-1)}{2}$. Since $|\text{Lines}(\mathbb{Z}_p^{n_0})| = \frac{p^{n_0}-1}{p-1} > \frac{p(p-1)}{2}$, for $n_0 \geq 3$, and every $\mathbf{q} \in \text{Lines}(\mathbb{Z}_p^{n_0})$ is contained in a distinct line through the origin, we may pick $\mathbf{q} \in \text{Lines}(\mathbb{Z}_p^n) \setminus \mathbb{Z}_p(PS - PS) \neq \emptyset$. Now let $B \in \mathbb{Z}_p^{n_0 \times (n_0-1)}$ denote a basis satisfying $\mathbf{q}^\perp = B \mathbb{Z}_p^{n_0-1}$. I claim that $|B^T PS| = |PS|$. Assume not, then there

exists distinct $\mathbf{x}, \mathbf{y} \in PS$ such that

$$\begin{aligned} B^T \mathbf{x} \equiv B^T \mathbf{y} \pmod{p\mathbb{Z}^{n_0-1}} &\Leftrightarrow B^T(\mathbf{x} - \mathbf{y}) \equiv \mathbf{0} \pmod{p\mathbb{Z}^{n_0-1}} \\ &\Leftrightarrow (\mathbf{x} - \mathbf{y}) \in (B\mathbb{Z}_p^{n_0-1})^\perp = \mathbb{Z}_p \mathbf{q}, \end{aligned}$$

which contradicts to our assumption on \mathbf{q} . Therefore, the algorithm is indeed guaranteed to find a valid projection as needed.

After the first for loop, we have construction $P \in \mathbb{Z}_p^{2 \times n}$ satisfying $|PS| = |S|$, where $|S| < p < 2|S| - 2$. By Lemma 5.5.7, there exists $\mathbf{a} \in \mathbb{Z}_p^2$ satisfying (1) and (2) for the set PS . Since (1) and (2) holds for any non-zero multiple of \mathbf{a} , i.e. any vector defining the same line as \mathbf{a} , we may restrict the search to elements of $\text{Lines}(\mathbb{Z}_p^2)$. Therefore, by trying all $\frac{p^2-1}{p-1} = p+1$ elements of $\text{Lines}(\mathbb{Z}_p^2)$ the algorithm is guaranteed to find a valid \mathbf{a} for the PS . Noting that $\langle \mathbf{a}, P\mathbf{y} \rangle \equiv \langle P^t \mathbf{a}, \mathbf{y} \rangle \pmod{p}$, we get that $P^t \mathbf{a}$ satisfies (1) and (2) for the set S , as needed.

Runtime: For $n = 1$, the runtime is constant. We assume $n \geq 2$. Here the first for loop is executed $n - 2$ times. For each loop iteration, we run through $\mathbf{q} \in \text{Lines}(\mathbb{Z}_p^{n_0})$, until we find one inducing a good projection matrix B . From the above analysis, we iterate through at most $\binom{|S|}{2} < \frac{p(p-1)}{2}$ elements $\mathbf{q} \in \text{Lines}(\mathbb{Z}_p^{n_0})$ before finding a good projection matrix. For each \mathbf{q} , we build a basis matrix for B for \mathbf{q}^\perp which can be done in using $\text{poly}(n, \log p)$ arithmetic operations via standard methods. Next we check for collisions against each pair $\mathbf{x}, \mathbf{y} \in PS$, which can be done using $O(|S|) = O(p)$ iterations over S . Therefore, at each loop iteration we enumerate over S at most p^3 times while performing only polynomial time computations. Hence, the total number of operations (excluding the time needed to output the elements of S) is at most $\text{poly}(n, \log p)p^4$.

For the last phase, we run through the elements in $\text{Lines}(\mathbb{Z}_p^2)$, where $|\text{Lines}(\mathbb{Z}_p^2)| = p+1$. The validity check for $\mathbf{a} \in \text{Lines}(\mathbb{Z}_p^2)$, simply requires computing both the quantities (1) and (2). To compute $|\{\mathbf{y} \in S : \langle \mathbf{y}, \mathbf{a} \rangle \equiv 0 \pmod{p}\}|$ we iterate once over the

set S and count how many zero dot products there are. To compute $|\{\langle \mathbf{a}, \mathbf{y} \rangle : \mathbf{y} \in S\}|$, we first iterate over all residues in \mathbb{Z}_p . Next for each residue $i \in \mathbb{Z}_p$, if we find $\mathbf{y} \in S$ satisfying $\langle \mathbf{a}, \mathbf{y} \rangle \equiv i \pmod{p}$, we increment our counter by one, and otherwise continue. Hence for any specific $\mathbf{a} \in \mathbb{Z}_p^2$, we iterate over the set S exactly $p + 1$ times, performing $\text{poly}(n, \log p)p^2$ operations. Hence, over the whole loop we perform $O(p^2)$ iterations of the set S , and perform $\text{poly}(n, \log p)p^3$ operations.

Therefore, over the whole algorithm we iterate over the set S at most np^3 times, and perform at most $\text{poly}(n, \log p)p^4$ operations. Furthermore, not counting the space needed to iterate over the set S , the space used by the the algorithm is $\text{poly}(n, \log p)$. \square

Using the Good-Vector algorithm, we give a completely deterministic construction for Lattice Sparsifiers.

Algorithm 5.8 Algorithm Lattice-Sparsifier(K, \mathcal{L}, t)

Input: $(\mathbf{0}, r, R)$ -centered convex body $K \subseteq \mathbb{R}^n$ with distance oracle D_K for $\|\cdot\|_K$, basis $B \in \mathbb{Q}^{n \times n}$ for \mathcal{L} , and $t \geq 0$.

Output: (K, t) sparsifier for \mathcal{L}

- 1: $K \leftarrow K \cap -K$
 - 2: Compute $\mathbf{y} \in \text{Shortest-Vectors}(K, \mathcal{L}, \frac{1}{3})$
 - 3: $\lambda \leftarrow D_K(\mathbf{y}, \frac{1}{3})$; $\epsilon \leftarrow 9^{-(n+3)}$
 - 4: $k \leftarrow \lfloor \ln(\frac{3}{4}\frac{t}{\lambda} + 1) / \ln 4 \rfloor$
 - 5: $\mathcal{L}_0 \leftarrow \mathcal{L}$; $B_0 \leftarrow B$
 - 6: **for** i **in** 0 **to** $k - 1$ **do**
 - 7: $S \leftarrow \text{Lattice-Enum}(4^i(1 - \epsilon)\lambda K, \mathcal{L}_i, \epsilon\lambda r)$
 - 8: Compute $N \leftarrow |S|$
 - 9: **if** $N > 3$ **then**
 - 10: Compute $B_i^* \leftarrow B_i^{-T}$, a basis for \mathcal{L}_i^*
 - 11: Compute prime p satisfying $N < p < 2N - 2$
 - 12: $\mathbf{a} \leftarrow \text{Good-Vector}(B_i^{*T}S \pmod{p\mathbb{Z}^n}, p)$
 - 13: Compute $\mathcal{L}_{i+1} \leftarrow \{\mathbf{y} \in \mathcal{L}_i : \langle B_i^* \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}$ and basis B_{i+1} for \mathcal{L}_{i+1}
 - 14: **else**
 - 15: $\mathcal{L}_{i+1} \leftarrow \mathcal{L}_i$; $B_{i+1} \leftarrow B_i$
 - 16: **return** \mathcal{L}_k
-

Theorem 5.5.10 (Correctness of Lattice Sparsifier). *Given a γ -symmetric convex*

body $K \subseteq \mathbb{R}^n$, Algorithm 5.8 returns a (K, t) sparsifier for \mathcal{L} satisfying $G(tK, \mathcal{L}) \leq 2^{O(n)}\gamma^{-n}$ using $2^{O(n)} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space.

Analysis of Lattice Sparsifier.

Correctness: We show that the outputted lattice is a (K, t) sparsifier for \mathcal{L} . By Lemma 5.5.2 it suffices to show that the algorithm outputs a $(K \cap -K, t)$ sparsifier, which justifies the switch in line 2 from K to $K \cap -K$. In what follows, we therefore assume that K is symmetric.

We first claim that $\lambda \leq 2\lambda_1(K, \mathcal{L})$. To see by the guarantee on $\text{Shortest-Vector}(K, \mathcal{L}, \frac{1}{3})$, we have that $\|y\|_K \leq (1 + \frac{1}{3})\lambda_1(K, \mathcal{L})$. Next, by the guarantee on D_K , we have that

$$\lambda = D_K(\mathbf{y}, \frac{1}{3}) \leq (1 + \frac{1}{3})\|\mathbf{y}\|_K \leq (1 + \frac{1}{3})^2\lambda_1(K, \mathcal{L}) \leq 2\lambda_1(K, \mathcal{L}),$$

as needed.

Claim 1: for each i , $0 \leq i \leq k$, we have that

$$(1) \quad \forall \mathbf{x} \in \mathbb{R}^n, d_K(\mathcal{L}_i, \mathbf{x}) \leq d_K(\mathcal{L}, \mathbf{x}) + \frac{4}{3}(4^i - 1)\lambda.$$

$$(2) \quad G(4^i\lambda, \mathcal{L}_i) = 9^{n+2}.$$

Proof. We establish the claim by induction on i . For $i = 0$, we have that $\mathcal{L}_0 = \mathcal{L}$. Therefore, \mathcal{L}_0 trivially satisfies property (1). Next, since $\lambda \leq 2\lambda_1(K, \mathcal{L})$, by Lemma 5.3.1 we have that $G(\lambda K, \mathcal{L}_0) \leq (2(2) + 1)^n = 5^n < 9^{n+2}$. Hence \mathcal{L}_0 also satisfies (2).

We now prove the claim for $i \geq 1$. Let S denote the set outputted by $\text{Lattice-Enum}(4^{i-1}(1 - \epsilon)\lambda K, \mathcal{L}_{i-1}, \epsilon\lambda r)$. By the guarantees on Lattice-Enum , the set S satisfies $4^{i-1}(1 - \epsilon)\lambda K \cap \mathcal{L}_{i-1} \subseteq S \subseteq (4^{i-1}(1 - \epsilon)\lambda K + \epsilon\lambda r B_2^n) \cap \mathcal{L}_{i-1}$. Since $r B_2^n \subseteq K$ and $i \geq 1$, we note that $4^{i-1}(1 - \epsilon)\lambda K + \epsilon\lambda r B_2^n \subseteq 4^i\lambda K$. Therefore, we have that

$$4^{i-1}(1 - \epsilon)\lambda K \cap \mathcal{L}_{i-1} \subseteq S \subseteq 4^{i-1}\lambda K \cap \mathcal{L}_{i-1} \tag{5.5.2}$$

Let N denote $|S|$ (line 8). By (5.5.2) and the induction hypothesis, we have that

$$|4^{i-1}(1 - \epsilon)\lambda K \cap \mathcal{L}_{i-1}| \leq N \leq |4^{i-1}\lambda K \cap \mathcal{L}_{i-1}| \leq G(4^{i-1}\lambda K, \mathcal{L}) \leq 9^{n+2}$$

Assume $N \leq 3$. Then the algorithm sets $\mathcal{L}_i = \mathcal{L}_{i-1}$ and $B_i = B_{i-1}$. By the induction hypothesis, for $\mathbf{x} \in \mathbb{R}^n$ we have that

$$d_K(\mathcal{L}_i, \mathbf{x}) = d_K(\mathcal{L}_{i-1}, \mathbf{x}) \leq d_K(\mathcal{L}, \mathbf{x}) + \frac{4}{3}(4^{i-1} - 1)\lambda \leq d_K(\mathcal{L}, \mathbf{x}) + \frac{4}{3}(4^i - 1)\lambda,$$

and hence \mathcal{L}_i satisfies (1). Next, by (5.5.2) we have that $|4^i(1 - \epsilon)\lambda K \cap \mathcal{L}_i| \leq N \leq 3$. Therefore by Lemma 5.4.1, we have that

$$\begin{aligned} G(4^{i+1}\lambda K, \mathcal{L}_{i+1}) &\leq (2 \cdot 4(1/(1 - \epsilon)) + 1)^n |4^i(1 - \epsilon)\lambda K \cap \mathcal{L}_{i+1}| \\ &\leq 9^n(1 + 2\epsilon)^n \cdot 3 \leq 9^{n+2}, \end{aligned}$$

where the last two inequalities follow since $\epsilon \leq 9^{-(n+3)}$. Therefore \mathcal{L}_i satisfies requirement (2) as needed.

Assume $N > 3$. Here we first compute $N < p < 2N - 2$ (which exists by Bertrand's Postulate), and a dual basis B_{i-1}^* for \mathcal{L}_{i-1}^* .

Claim 2: $|B_{i-1}^{*T} S \pmod{p\mathbb{Z}^n}| = N$

Proof. Since $|S| = N$, if the claim is false, there exists distinct $\mathbf{x}, \mathbf{y} \in \mathcal{L}$ such that

$$B_{i-1}^{*T}\mathbf{x} \equiv B_{i-1}^{*T}\mathbf{y} \pmod{p\mathbb{Z}^n} \Leftrightarrow B_{i-1}^{*T}(\mathbf{x} - \mathbf{y}) \equiv \mathbf{0} \pmod{p\mathbb{Z}^n} \Leftrightarrow \mathbf{x} - \mathbf{y} \in p\mathcal{L}_{i-1}.$$

Since $\mathbf{x}, \mathbf{y} \in 4^{i-1}\lambda K$ and K is symmetric, we have that $\mathbf{x} - \mathbf{y} \in 2 \cdot 4^{i-1}K \cap p\mathcal{L}_{i-1}$. Let $\mathbf{z} = \mathbf{x} - \mathbf{y} \in p\mathcal{L}_{i-1}$. We examine the vector $s\frac{\mathbf{z}}{p}$ for $s \in \mathbb{Z}$ satisfying $|s| \leq \lfloor \frac{p}{2} \rfloor = \frac{p-1}{2}$ (since p is odd). Since $\frac{\mathbf{z}}{p} \in \mathcal{L}_{i-1}$, we have that $s\frac{\mathbf{z}}{p} \in \mathcal{L}_{i-1}$ and

$$\begin{aligned} s\frac{\mathbf{z}}{p} \in \left| \frac{s}{p} \right| \cdot 2 \cdot 4^{i-1}K &\subseteq \left(\frac{p-1}{2p} \right) 2 \cdot 4^{i-1}K = \left(1 - \frac{1}{p} \right) 4^{i-1}K \\ &\subseteq (1 - \epsilon)4^{i-1}K \end{aligned}$$

where the last inequality follows since $p < 2N - 2 \leq 2 \cdot 9^{n+2}$ and $\epsilon = 9^{-(n+3)}$. Then, since s can take on $2\lfloor \frac{p}{2} \rfloor + 1 = p$ different values, the set $(1 - \epsilon)4^{i-1}K$ contains at least p lattice points in \mathcal{L}_{i-1} . However, by construction of N , we have that $|(1 - \epsilon)4^{i-1}K \cap \mathcal{L}_{i-1}| \leq N < p$, a clear contradiction. The claim thus holds. \square

Next, the algorithm computes $\mathbf{a} \leftarrow \text{Good-Vector}(B_i^{*T}S \pmod{p\mathbb{Z}^n}, p)$, and sets $\mathcal{L}_i = \{\mathbf{y} \in \mathcal{L} : \langle B^* \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}$. From Claim 2, equation 5.5.2, and the guarantees on Good-Vector, we get that

- (a) $|4^{i-1}(1-\epsilon)\lambda K \cap \mathcal{L}_i| = |\{\mathbf{y} \in 4^{i-1}(1-\epsilon)\lambda K \cap \mathcal{L}_{i-1} : \langle B^* \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}| \leq 6$.
- (b) $|\{\langle B^* \mathbf{a}, \mathbf{y} \rangle \pmod{p} : \mathbf{y} \in 4^{i-1}\lambda K \cap \mathcal{L}_{i-1}\}| \geq \frac{p}{4} + 1$.

From here, using the identical analysis as in Theorem 5.5.8, from (a) above we get that $\forall \mathbf{x} \in \mathbb{R}^n$, $d_K(\mathcal{L}_i, \mathbf{x}) \leq d_K(\mathcal{L}_{i-1}, \mathbf{x}) + 4^i \lambda$. Now by the induction hypothesis on \mathcal{L}_{i-1} , we get that

$$d_K(\mathcal{L}_{i-1}, \mathbf{x}) + 4^i \lambda \leq d_K(\mathcal{L}, \mathbf{x}) + \frac{4}{3}(4^{i-1} - 1) + 4^i \lambda = d_K(\mathcal{L}, \mathbf{x}) + \frac{4}{3}(4^i - 1)\lambda.$$

Therefore \mathcal{L}_i satisfies (1) as needed. Using (b) and Lemma 5.4.1, we have that

$$\begin{aligned} G(4^i \lambda K, \mathcal{L}_i) &\leq (2 \cdot 4 \cdot (1/(1-\epsilon)) + 1)^n |4^{i-1}(1-\epsilon)\lambda K \cap \mathcal{L}_i| \\ &\leq 9(1+2\epsilon)^n \cdot 6 \leq 9^{n+2}. \end{aligned}$$

Therefore \mathcal{L}_i satisfies (2) as needed. The claim thus follows. \square

Given Claim 1, we will show that \mathcal{L}_k is (K, t) sparsifier for \mathcal{L} . By our choice of k , note that $\frac{4}{3}(4^k - 1)\lambda \leq t \leq 4 \cdot \frac{4}{3}(4^{k+1} - 1)\lambda$. By the claim, for $\mathbf{x} \in \mathbb{R}^n$, $d_K(\mathcal{L}_k, \mathbf{x}) \leq d_K(\mathcal{L}, \mathbf{x}) + \frac{4}{3}(4^k - 1)\lambda \leq d_K(\mathcal{L}, \mathbf{x}) + t$. It therefore only remains to bound $G(tK, \mathcal{L}_k)$. By the previous bounds, note that

$$\frac{t}{4^k \lambda} \leq \frac{4(4^{k+1} - 1)\lambda}{3 \cdot 4^k \lambda} \leq \frac{16}{3} \leq 6$$

Therefore by the claim and Lemma 5.4.1, we have that

$$G(tK, \mathcal{L}_k) \leq (2 \cdot 6 + 1)^n G(4^k \lambda K, \mathcal{L}_k) \leq 13^n \cdot 9^{n+2} = 2^{O(n)}$$

as needed. The algorithm therefore returns a valid (K, t) sparsifier for \mathcal{L} as needed.

Runtime: The algorithm first runs the Shortest-Vectors on K and \mathcal{L} , which takes $2^{O(n)}$ $\text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space. Next, the for loop on line 6 iterates $k = \lfloor \ln(\frac{3}{4}\frac{t}{\lambda} + 1) / \ln 4 \rfloor = \text{poly}(\cdot)$ times.

Each for loop iteration, indexed by i satisfying $0 \leq i \leq k - 1$, consists of computations over the set $S \leftarrow \text{Lattice-Enum}(4^i(1 - \epsilon)\lambda K, \mathcal{L}_i, \epsilon\lambda r)$. For the intended implementation, we do not store the set S explicitly. Every time the algorithm needs to iterate over the S , we implement this by performing a call to $\text{Lattice-Enum}(4^i(1 - \epsilon)\lambda K, \mathcal{L}_i, \epsilon\lambda r)$. Furthermore, note the algorithm only interacts with S by iterating over its elements, and hence the implemented interface suffices. Now at the loop iteration indexed by i , we do as follows:

- (1) Compute $N = |S|$. This is implemented by iterating over the elements of S and counting, and so by the guarantees of Lattice-Enum requires at most $2^{O(n)}G(4^i\lambda K, \mathcal{L}_i) \text{poly}(\cdot) = 2^{O(n)} \text{poly}(\cdot)$ time (by Claim 1) and $2^n \text{poly}(\cdot)$ space.
- (2) If $N \leq 3$, we keep the same lattice and skip to the next loop iteration. If $N > 3$, continue.
- (3) Compute $B_i^* = B_i^{-T}$. This can be done in $\text{poly}(\cdot)$ time and space.
- (4) Compute prime p satisfying $N < p < 2N - 2$. Such a prime can be computed by trying all integers in the previous range and using trial division. This takes at most $O(N^2 \text{poly}(\log N)) = 2^{O(n)}$ time and $\text{poly}(n)$ space.
- (5) Call $\text{Good-Vector}(B^{T*}S \pmod{p\mathbb{Z}^n}, p)$. By the guarantees on Good-Vector, the algorithm performs $\text{poly}(n, \log p)p^4 = 2^{O(n)}$ operations and iterates at most $np^3 = 2^{O(n)}$ times over the set $B^{T*}S \pmod{p\mathbb{Z}^n}$. These iterations can be performed $2^{O(n)} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space by the guarantees on Lattice-Enum.

(6) Compute a basis B_{i+1} for the new lattice $\mathcal{L}_{i+1} = \{\mathbf{y} \in \mathcal{L}_i : \langle B^{*T} \mathbf{a}, \mathbf{y} \rangle \equiv 0 \pmod{p}\}$.

This can be done in $\text{poly}(\cdot)$ time using standard methods.

From the above analysis, we see that the entire algorithm runs in $2^{O(n)} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space as needed. \square

5.6 Finding a Central Lattice Point

In this section, we give a Las Vegas algorithm to find a lattice point near an approximate center of mass of a convex body K . The algorithm presented here will have applications to integer programming, which we elaborate on in Chapter 7.

We present the algorithm below.

Algorithm 5.9 Central-Lat-Pt(K, \mathcal{L}, ϵ)

Input: Weak membership oracle O_K for a (\mathbf{a}_0, r, R) -centered convex body $K \subseteq \mathbb{R}^n$, a basis $B \in \mathbb{Q}^{n \times n}$ for \mathcal{L} , and $\epsilon \in (0, 1)$.

Output: Vector $\mathbf{b} \in K$ and $\mathbf{y} \in \mathcal{L}$ satisfying the conditions of Theorem 5.6.1.

- 1: Let $\mathbf{b} \leftarrow \text{Approx-Mass-Center}(K)$
 - 2: Pick $\mathbf{y} \in \text{Approx-Closest-Vectors}(K - \mathbf{b}, \mathcal{L}, \mathbf{b}, \epsilon)$
 - 3: **return** (\mathbf{b}, \mathbf{y})
-

Theorem 5.6.1 (Correctness of Central-Lat-Pt). *On input K, \mathcal{L} and $\epsilon > 0$ as above, in expected $2^{O(n)}(1 + \frac{1}{\epsilon})^n \text{poly}(\cdot)$ time, using $2^n \text{poly}(\cdot)$ space and $\text{poly}(\cdot)$ randomness, algorithm 5.9 outputs a vector $\mathbf{b} \in K$ and vector $\mathbf{y} \in \mathcal{L}$ such that*

- (1) K is $(\mathbf{b}, \frac{r}{2\sqrt{n(n+1)}}, 2R)$ -centered and $K - \mathbf{b}$ is $\frac{1}{5}$ -symmetric.
- (2) $\|\mathbf{y} - \mathbf{b}\|_{K-\mathbf{b}} \leq (1 + \epsilon)d_{K-\mathbf{b}}(\mathcal{L}, \mathbf{b})$.

Analysis of Approx-Lat-Cont.

Correctness: The correctness of the algorithm follows directly by the guarantees on algorithms Approx-Mass-Center (see algorithm 4.6) and Approx-Closest-Vectors (see algorithm 5.6).

Runtime: The call to Approx-Mass-Center requires expected $2^{O(n)}$ poly(\cdot)-time and poly(\cdot) space. Since $K - \mathbf{b}$ is $\frac{1}{5}$ -symmetric, the call of to Approx-Closest-Vectors requires $2^{O(n)}(1 + \frac{1}{\epsilon})^n$ time and 2^n poly(\cdot) space. The desired runtime bound thus follows. Since the only randomness comes from Approx-Mass-Center, where the amount of randomness used is polynomial, we get that the entire algorithm uses only polynomial randomness, as needed. \square

Acknowledgments. We would like to express our gratitude to Matthias Köppe for suggesting the use of reverse search to save space in our covering and enumeration algorithms, to Gabor Kun for suggesting the dimension reduction technique used in derandomizing the lattice sparsifier construction, and to Daniele Micciancio and Panagiotis Voulgaris for useful discussions relating to the research in this chapter.

5.7 Conclusion

The study of lattice problems, such as the SVP and CVP, has lead to many fundamental discoveries in computer science and the geometry of numbers. In this chapter, we have performed an in-depth study of these lattices problems under general norms. For our main contributions, we gave efficient deterministic algorithms for the SVP, CVP and $(1 + \epsilon)$ -CVP under general norms, that yield the only known deterministic alternatives to the previous AKS sieve based algorithms for these problems. Though our running times are generally $2^{O(n)}$ factors larger than the AKS alternatives, our algorithms save considerably on space compared to AKS. In particular, our $(1 + \epsilon)$ -CVP algorithm requires only 2^n space as opposed to $(1 + 1/\epsilon)^n$ space required by the AKS sieve approaches. Furthermore, the deterministic guarantees on our algorithms will be helpful for their application to Integer Programming, which we describe in Chapter 7.

From a more general perspective, we have introduced new geometric techniques to

the study of lattice problems. In particular, we show the usefulness of ellipsoid coverings and M-Ellipsoids to lattice algorithms, and present a new type of complexity guarantee for lattice point enumeration (i.e. the $G(K, \mathcal{L})$ bound). Lastly, our lattice sparsifier construction for general norms yields, to the best of our knowledge, a new structural result for lattices. We note that the core idea behind the lattice sparsifier, i.e. the random sublattice restriction, has appeared previously in the context of proving NP-hardness for SVP under ℓ_p norms [75, 62]. Furthermore, it remains an outstanding open problem to give a deterministic NP-hardness reduction for SVP under the ℓ_2 norm. We are therefore hopeful that our analysis of the lattice sparsifier construction and its subsequent derandomization may find applications beyond what has been described here.

Future Research. A first important problem we would like to explore is that of finding a polynomial space and $2^{O(n)}$ time algorithm for the CVP under ℓ_2 . As shown in Theorem 5.1.10, such an algorithm would immediately reduce the space complexity of all our algorithms to polynomial. As a first step towards such an algorithm, we would like to explore the problem of finding a succinct implicit representation of the voronoi relevant vectors. One question in this vein is as follows: for an n dimensional lattice \mathcal{L} is there a collection of vectors $v_1, \dots, v_{cn} \in \mathbb{R}^n$, such that the vectors in VR are expressible as integer combinations of v_1, \dots, v_{cn} of norm at most $C\sqrt{n}$ (of the coefficient vector), for $c, C \geq 1$ absolute constants? The existence of such a representation, would immediately imply a $2^{O(n)}$ time and $\text{poly}(n)$ space algorithm for CVP with preprocessing (i.e. where the list is given as advice), and hence yield important step towards this problem.

A next question pertains to the current structure of our algorithms. At a high level, we show that the SVP, CVP and $(1 + \epsilon)$ -CVP under general norms all reduce to ℓ_2 CVP. The reduction however, which passes through the M-Ellipsoid computation,

ellipsoid coverings (generated by a parallelepiped tiling), and lattice point enumeration inside an ellipsoid, is quite complex and very time inefficient (though quite space efficient). A natural question is therefore whether all this complexity is needed. In particular, is there a direct way of achieving the guarantees of our algorithms without passing through this complex reduction to ℓ_2 ?

Lastly, we would like to explore computational lower bounds for the SVP and CVP. Even though such bounds are in general dependent on $P \neq NP$, we note that all our algorithms require only query access to the norm in question, and hence are amenable to information theoretic lower bounds. Here it would be interesting to establish that our algorithms for SVP and $(1 + \epsilon)$ -CVP algorithms are nearly optimal in this model, in the same vein as the volume estimation algorithm from Chapter 4.

CHAPTER VI

GEOMETRY OF THE DISCRETE GAUSSIAN AND THE FLATNESS THEOREM

In this chapter, we give a tighter proof of Kinchine’s flatness theorem in the geometry of numbers. To prove this result we develop a tight characterization of the smoothing parameter of a lattice, as well as nearly matching upper and lower bounds for norm expectations of discrete Gaussian random variables. The work in this chapter is based on parts of the paper [27] (joint with Kai-Min Chung, Feng Hao Liu, and Chris Peikert).

6.1 Introduction

One of the most famous results in the Geometry of Numbers is Minkowski’s theorem. Minkowski answered the question “when does a centrally symmetric convex body contain a non-zero integer point?”. He showed that any symmetric convex body with large enough volume indeed contains such a point. His discovery laid the way for many significant advances in number theory, and has had ramifications in many other areas including combinatorics, computational complexity, and cryptography.

A natural extension to the above question is “when does a convex body contain an integer point?”. Here the situation is very different from the above since large volume no longer yields a sufficient condition. One may examine the example of a band in between two consecutive integral hyperplanes, to see that a convex set may have infinite volume and yet not contain any integers. Instead of volume, what turns out to be true is that any convex set that is not too “flat” must contain an integer point.

This result, known as Kinchine’s flatness theorem is a powerful tool in the geometry of numbers. Stated precisely, the flatness theorem says that for any convex body K in \mathbb{R}^n , either $\mu(K, \mathbb{Z}^n) = \inf\{s \geq 0 : \mathbb{Z}^n + sK = \mathbb{R}^n\} \leq 1$ (covering radius is small) or

$$\lambda_1((K - K)^*, \mathbb{Z}^n) = \inf_{y \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \text{width}_K(y) \leq f(n) \text{ (width function of } K)$$

K has “integer width” bounded by a function of dimension. We note that the above statement generalizes to any n dimensional lattice (which we state below), and is not restricted to the integers. To understand the relation to the original question, note that if K has covering radius less than 1, then not only does K contain an integer point, but so does every translation of K , i.e. a seemingly much stronger statement.

A major application of the flatness theorem is to the Integer Programming Problem (IP), i.e. the problem of deciding whether a convex set contains an integer point. Here, the idea is that if a convex body K has integer width $> f(n)$, then we can immediately correctly decide that $K \cap \mathbb{Z}^n \neq \emptyset$. Otherwise if the integer width is $\leq f(n)$, then we can use a short integer width direction to decompose the feasibility problem into subproblems along at most $f(n)$ integral hyperplanes. Most of the known algorithms for IP (achieving any reasonable complexity guarantee) rely on this type of approach (see for example [84, 70, 65, 36]). For more details on the application of the flatness theorem to IP, we refer the interested reader to Chapter 7.

Due to its importance to IP and other problems, many works have been devoted to the task of improving the bounds on the flatness function [76, 7, 80, 67, 10, 12, 13]. The main result of this chapter is a quantitative improvement on a theorem of Banaszczyk [10], which reduces the task of bounding the flatness constant for dimension n to bounding the ℓ^* estimate in Convex Geometry (see section 4.4.1 for a thorough discussion of ℓ^*). Our main result is stated as follows:

Theorem 6.1.1. *For a convex body K and lattice \mathcal{L} in \mathbb{R}^n , we have that*

$$1 \leq \mu(K, \mathcal{L}) \lambda_1((K - K)^*, \mathcal{L}^*) \leq \frac{8\sqrt{2}}{\pi} \ell\ell^*(K) \quad (6.1.1)$$

where

$$\ell\ell^*(K) \stackrel{\text{def}}{=} \inf_{\substack{A \in \mathbb{R}^{n \times n}, \det(T)=1 \\ \mathbf{x} \in K}} \ell_{K-\mathbf{x}}(A) \ell_{(K-\mathbf{x})^*}(A^{-T})$$

We note that definition of $\ell\ell^*(K)$ is slightly more general than the one presented in Section 4.4.1. In section 4.4.1 we only needed to define $\ell\ell^*$ for symmetric convex bodies, whereas here we require it for general bodies. To compensate for the lack of symmetry, we allow for recentering the body K . For $K = B_2^n$, $\ell\ell^*(B_2^n) = n$. For a general symmetric convex body K , it was shown by Pisier[106] that $\ell\ell^*(K) = 4n(1 + \frac{1}{2} \log_2 n)$. For asymmetric bodies, the estimates are significantly worse, where the best known bound is $O(n^{\frac{4}{3}} \text{polylog}(n))$ due to Rudelson [113]. It remains a major open problem in Convex geometry to show that $\ell\ell^*(K) = O(n \log n)$ for asymmetric bodies (which is tight for $K = B_\infty^n$).

For the above theorem, we note that the lower bound $1 \leq \mu(K, \mathcal{L}) \lambda_1((K - K)^*, \mathcal{L}^*)$ is classical (see Lemma 2.4.6), and hence our focus in this Chapter is dedicated to the upper bound.

The above theorem, which yields the current best asymptotic upper bounds on the flatness function $f(n)$, was first proved by Banaszczyk in [12] with very large (and hard to compute) hidden constants. More precisely, to prove 6.1.1 Banaszczyk relies on Talagrand's majorizing measure theorem. Here the majorizing measure theorem is used to upper bound expectations of the form $E[\|Y\|_K]$ for a sub-Gaussian random variable Y (where Y represents a certain discrete distribution over \mathcal{L}) by the corresponding Gaussian expectation. However, the constants in this reduction are to the best of the author's knowledge, unknown and presumed to be quite large. In this chapter, we follow same outline as Banaszczyk's proof while avoiding the use of majorizing measure theorem. At a high level, we show that the generality of the

majorizing measure theorem is unnecessary by relying on the lattice structure more directly. Furthermore, we provide new structural insights for lattices which may be useful elsewhere.

We remark that even with the best possible bound on $\ell\ell^*$ (i.e. $O(n \log n)$), there would still be a gap with respect to the best known lower bound for $f(n)$, i.e. $f(n) = \Omega(n)$ (see [10]). Indeed, it is conjectured that the $f(n) = \Theta(n)$. However, resolving this conjecture would seem to require completely different techniques than the ones described here.

6.2 Preliminaries

Discrete Gaussian measures. We define the Gaussian function $\rho_{n,s,\mathbf{c}} : \mathbb{R}^n \rightarrow \mathbb{R}_+$, $n \in \mathbb{N}$, $s > 0$, and $\mathbf{c} \in \mathbb{R}^n$ by $\rho_{n,s,\mathbf{c}}(\mathbf{x}) = e^{-\pi\|\frac{\mathbf{x}-\mathbf{c}}{s}\|_2^2}$ for $\mathbf{x} \in \mathbb{R}^n$. We define $\rho_{n,s} \stackrel{\text{def}}{=} \rho_{n,s,\mathbf{0}}$. For a countable subset $T \subseteq \mathbb{R}^n$, we define $\rho_{n,s,\mathbf{c}}(T) = \sum_{\mathbf{x} \in T} \rho_{n,s,\mathbf{c}}(\mathbf{x})$. When the context is clear, we shall often drop the n in the notation, and write ρ_s , $\rho_{s,\mathbf{c}}$ for $\rho_{n,s}$ and $\rho_{n,s,\mathbf{c}}$.

For any countable subset $T \subseteq \mathbb{R}^n$ for which $\rho_{s,\mathbf{c}}(T)$ converges, define the discrete Gaussian distribution $D_{T,s,\mathbf{c}}$ over T by

$$D_{T,s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(T)} \quad \forall \mathbf{x} \in T.$$

In this chapter, we will examine the discrete Gaussian over a lattice \mathcal{L} or one of its cosets, i.e., where $T = \mathcal{L} + \mathbf{c}$, for some $\mathbf{c} \in \mathbb{R}^n$. (In all this case, $\rho_{s,\mathbf{c}}(T)$ converges.)

Recalling the notation of section 2.1.3, $D_{n,s,\mathbf{c}}$ denotes the continuous Gaussian distribution over \mathbb{R}^n , and that $\gamma_{n,s}$ denotes the n dimensional Gaussian measure. We call $X \in \mathbb{R}^n$ a standard n -dimensional Gaussian if X is distributed as $D_{n,\sqrt{2\pi}}$. For $K \subseteq \mathbb{R}^n$ a $\mathbf{0}$ -centered convex body, $p \geq 1$, and $A \in \mathbb{R}^{n \times n}$, we recall the definition of the ℓ -norm:

$$\ell_K^p(A) = \mathbb{E}[\|AX\|_K^p]^{\frac{1}{p}},$$

where X is standard Gaussian. Note that this is slightly more general than the definition in 4.4.1, since we allow varying $p \geq 1$. Here we write $\ell_K^2 \equiv \ell_K$.

Poisson Summation Formula. For an integrable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the Fourier transform $\hat{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ of f is

$$\hat{f}(\mathbf{y}) = \int_{\mathbb{R}^n} e^{-2\pi i \langle \mathbf{y}, \mathbf{x} \rangle} f(\mathbf{x}) d\mathbf{x}, \quad \mathbf{y} \in \mathbb{R}^n.$$

For $f(\mathbf{x}) = \rho_{n,s,\mathbf{c}}(\mathbf{x})$, from classical analysis we have that

$$\hat{\rho}_{n,s,\mathbf{c}}(\mathbf{x}) = s^n e^{-2\pi i \langle \mathbf{c}, \mathbf{x} \rangle} \rho_{n,\frac{1}{s}}(\mathbf{x}) \tag{6.2.1}$$

Let \mathcal{L} denote an n -dimensional lattice. For $f : \mathbb{R}^n \rightarrow \mathbb{R}$ “nice-enough”, the Poisson summation formula states that

$$\sum_{\mathbf{y} \in \mathcal{L}} f(\mathbf{y}) = \frac{1}{\det(\mathcal{L})} \sum_{\mathbf{y} \in \mathcal{L}^*} \hat{f}(\mathbf{y}) \tag{6.2.2}$$

As for the niceness condition, the above formula holds in particular if f is continuous, integrable, and satisfies $|f(\mathbf{x})| + |\hat{f}(\mathbf{x})| \leq C(1 + \|\mathbf{x}\|_2)^{-n-\delta}$, for some constants $C, \delta > 0$. In this chapter, we will only apply the formula to $\rho_{n,s,\mathbf{c}}$, where the above conditions hold.

For more information on the Fourier transform and its properties, the reader may consult [103].

The smoothing parameter. We recall the definition of the smoothing parameter from [90].

Definition 6.2.1 (Smoothing Parameter). For a lattice \mathcal{L} and real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\mathcal{L})$ is the smallest $s > 0$ such that $\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \epsilon$.

At a high level, the smoothing parameter helps us measure how close the discrete Gaussian distribution from the continuous Gaussian. Making this intuition quantitative will be the crucial task of the rest of the chapter.

6.3 Bounding the Smoothing Parameter

The following theorem gives a general bound on the smoothing parameter of \mathcal{L} with respect to any $\epsilon > 0$ and norm. The work of Banaszczyk [12] gives the first bounds of this type. His proof however passes through Talagrand’s majorizing measure theorem, which yields bounds with very large hidden constants. Here we avoid the use of majorizing measures by deriving our general norm bounds via comparison to the “optimal” norm for the lattice, i.e. the norm induced by the voronoi cell. Furthermore, our technique allows us to bounds with respect to arbitrary Gaussian moments with small explicit constants.

Theorem 6.3.1. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a lattice and $K \subseteq \mathbb{R}^n$ be a symmetric convex body. Let $X \in \mathbb{R}^n$ denote the n dimensional Gaussian with distribution γ . Then for any $\epsilon > 0$ and $p > 0$, we have that*

$$\eta_\epsilon(\mathcal{L}) \leq 2 \left(1 + \frac{1}{\epsilon}\right)^{\frac{1}{p}} \frac{\mathbb{E}[\|X\|_K^p]^{\frac{1}{p}}}{\lambda_1(K, \mathcal{L}^*)}$$

The following theorem shows that for constant ϵ , the smoothing parameter is essentially determined (up to constant factors) by the expected norm of the Gaussian under the norm induced by the voronoi cell.

Theorem 6.3.2. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be an n -dimensional lattice, $\mathcal{V} = \mathcal{V}(\mathcal{L}^*)$, and $X \in \mathbb{R}^n$ the Gaussian with distribution $D_{1,0}$. Then*

$$\frac{1}{3} \mathbb{E}[\|X\|_{\mathcal{V}}] \leq \eta_{\frac{1}{2}}(\mathcal{L}) \leq 3 \mathbb{E}[\|X\|_{\mathcal{V}}]$$

To prove the above theorems, we will need the following technical lemmas. We give their proofs and continue to the proofs of the main subsection theorems.

The following is a standard Gaussian tailbound for both continuous and discrete Gaussian (see Lemma 2.4 in [11] for the discrete Gaussian bound).

Lemma 6.3.3. *Let $X \in \mathbb{R}^n$ be distributed as either D_s or $D_{\mathcal{L},s}$ for a lattice $\mathcal{L} \subset \mathbb{R}^n$. For any $\mathbf{v} \in \mathbb{R}^n$ and $t > 0$, we have*

$$\Pr[\langle X, \mathbf{v} \rangle \geq t\|\mathbf{v}\|] \leq e^{-\pi(t/s)^2},$$

and for $\epsilon > 0$ we have

$$\Pr[\|X\|^2 \geq (1 + \epsilon)s^2 \frac{n}{2\pi}] \leq ((1 + \epsilon)e^{-\epsilon})^{n/2},$$

which for $0 < \epsilon < \frac{1}{2}$ is bounded by $e^{-n\epsilon^2/6}$.

The following well known lemma bounds how quickly the Gaussian measure of a symmetric set can decrease as it is shifted away from the origin. We include its proof for completeness.

Lemma 6.3.4. *Let $S \subseteq \mathbb{R}^n$ be symmetric (i.e., $S = -S$) measurable set. Then for any $\mathbf{y} \in \mathbb{R}^n$,*

$$\gamma_s(S + \mathbf{y}) \geq \gamma_s(S) \cdot \rho_s(\mathbf{y}).$$

Proof. By scaling S and \mathbf{y} , it suffices to prove the claim for $s = 1$. For any $t \in \mathbb{R}$, note that $\cosh(t) = \frac{1}{2}(e^t + e^{-t}) \geq 1$. We have

$$\begin{aligned} \gamma(S + \mathbf{y}) &= \int_S e^{-\pi\|\mathbf{y}-\mathbf{x}\|^2} d\mathbf{x} = \int_S \frac{1}{2}(e^{-\pi\|\mathbf{y}-\mathbf{x}\|^2} + e^{-\pi\|\mathbf{y}+\mathbf{x}\|^2}) d\mathbf{x} && \text{(symmetry of } S) \\ &= e^{-\pi\|\mathbf{y}\|^2} \int_S e^{-\pi\|\mathbf{x}\|^2} \cdot \frac{1}{2}(e^{2\pi\langle \mathbf{x}, \mathbf{y} \rangle} + e^{-2\pi\langle \mathbf{x}, \mathbf{y} \rangle}) d\mathbf{x} && \text{(expanding the squares)} \\ &\geq \rho(\mathbf{y}) \int_S \rho(\mathbf{x}) d\mathbf{x} = \rho(\mathbf{y}) \cdot \gamma(S). && \square \end{aligned}$$

The following simple lemma establishes a tight relationship between discrete Gaussian sums on \mathcal{L} and the Gaussian mass of the Voronoi cell. It will play a crucial role in the segway.

Lemma 6.3.5. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a lattice with Voronoi cell $\mathcal{V} = \mathcal{V}(\mathcal{L})$, and let $s > 0$. Then*

$$\frac{\rho_s(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho_s(\mathcal{L})} \leq 1 - \gamma_s(\mathcal{V}) \leq \rho_{2s}(\mathcal{L} \setminus \{\mathbf{0}\}).$$

Proof. By scaling \mathcal{L} , it suffices to prove the claim for $s = 1$. We first show the upper bound. Let $X \in \mathbb{R}^n$ be distributed as D_1 , and note that $1 - \gamma(\mathcal{V}) = \Pr[X \notin \mathcal{V}]$. By the union bound and Lemma 6.3.3,

$$\begin{aligned} \Pr[X \notin \mathcal{V}] &= \Pr\left[\bigcup_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \{\langle X, \mathbf{y} \rangle > \tfrac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle\}\right] \leq \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \Pr[\langle X, \mathbf{y} \rangle > \tfrac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle] \\ &\leq \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} e^{-\pi \|\mathbf{y}/2\|^2} = \rho_2(\mathcal{L} \setminus \{\mathbf{0}\}). \end{aligned}$$

We now prove the lower bound. Since \mathcal{V} tiles space with respect to L , by applying Lemma 6.3.4 with $S = \mathcal{V}$, we have

$$1 - \gamma(\mathcal{V}) = \gamma(\mathbb{R}^n \setminus \mathcal{V}) = \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \gamma(\mathcal{V} + \mathbf{y}) \geq \gamma(\mathcal{V}) \cdot \rho(\mathcal{L} \setminus \{\mathbf{0}\}).$$

Since for $x, y \geq 0$, we have that $1 - x \geq xy \Leftrightarrow 1 - x \geq \frac{y}{y+1}$, from the above we get that

$$1 - \gamma(\mathcal{V}) \geq \frac{\rho(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho(\mathcal{L} \setminus \{\mathbf{0}\}) + 1} = \frac{\rho(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho(\mathcal{L})}$$

as desired. □

The following lemma shows the optimality of the voronoi cell with respect to certain Gaussian expectations.

Lemma 6.3.6. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a lattice, and let $K \subseteq \mathbb{R}^n$ be a symmetric convex body satisfying $\lambda_1(K, L) \geq 2$. Then for any differentiable and non-decreasing function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, we have that*

$$\mathbb{E}[f(\|X\|_{\mathcal{V}})] \leq \mathbb{E}[f(\|X\|_K)]$$

where X is distributed as $D_{s, \mathbf{0}}$, $s > 0$, and $\mathcal{V} = \mathcal{V}(\mathcal{L})$.

Proof. Letting $\|\cdot\|$ denote an arbitrary norm, we have that

$$\begin{aligned} \mathbb{E}[f(\|X\|)] &= \int_{\mathbb{R}^n} f(\|\mathbf{x}\|) d\gamma_s(\mathbf{x}) = \int_{\mathbb{R}^n} \left(f(0) + \int_0^{\|\mathbf{x}\|} f'(t) dt \right) d\gamma_s(\mathbf{x}) \\ &= f(0) + \int_{\mathbb{R}^n} \int_0^{\|\mathbf{x}\|} f'(t) dt d\gamma_s(\mathbf{x}) = f(0) + \int_0^\infty f'(t) \gamma_s(\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| \geq t\}) dt \end{aligned}$$

Given the above identity and that $f' \geq 0$ (since f is non-decreasing), it now suffices to prove that for all $t > 0$

$$\begin{aligned} \gamma_s(\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{\mathcal{V}} \geq t\}) &\leq \gamma_s(\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_K \geq t\}) \Leftrightarrow \\ 1 - \gamma_s(\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{\mathcal{V}} \leq t\}) &\leq 1 - \gamma_s(\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_K \leq t\}) \Leftrightarrow \gamma_s(t\mathcal{V}) \geq \gamma_s(tK) \end{aligned}$$

Define $K' = \text{int}(K)$, i.e. $K' = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_K < 1\}$. Note that for $t > 0$, $(tK)' = tK'$.

Claim 1: For any $t > 0$, there exists an injective measure preserving (Lebesgue measure) map $T : tK' \rightarrow t\mathcal{V}$ satisfying that $\|T(\mathbf{x})\|_2 \leq \|\mathbf{x}\|_2 \forall \mathbf{x} \in K'$.

We first show the claim for $t = 1$. Define the map $c : \mathbb{R}^n \rightarrow \mathcal{L}$ which sends $\mathbf{x} \in \mathbb{R}^n$ to the lexicographically minimal (using the standard lexicographic ordering on the coordinates) lattice vector $\mathbf{y} \in L$ which is closest to \mathbf{x} under the ℓ_2 norm. Let the map $T : K' \rightarrow \mathcal{V}$ be defined by $T(\mathbf{x}) = \mathbf{x} - c(\mathbf{x})$. To see that the map is well defined, note that for any $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{x} - c(\mathbf{x})$ is closer to $\mathbf{0}$ than any other point in \mathcal{L} (since otherwise $c(\mathbf{x})$ would not be a closest lattice vector to \mathbf{x}), and hence $\mathbf{x} - c(\mathbf{x}) \in \mathcal{V}$ as needed. By definition $\|\mathbf{x} - c(\mathbf{x})\|_2 = \inf_{\mathbf{y} \in \mathcal{L}} \|\mathbf{x} - \mathbf{y}\|_2$ and hence $\|T(\mathbf{x})\|_2 \leq \|\mathbf{x}\|_2$ as needed. We now show that T is injective on K' . Assume not, i.e. there exists distinct $\mathbf{x}, \mathbf{y} \in K'$ such that $\mathbf{x} - c(\mathbf{x}) = \mathbf{y} - c(\mathbf{y})$. If this is the case then note that $\mathbf{0} \neq \mathbf{x} - \mathbf{y} = c(\mathbf{x}) - c(\mathbf{y}) \in \mathcal{L}$. From here we have that

$$\|\mathbf{x} - \mathbf{y}\|_K \leq \|\mathbf{x}\|_K + \|\mathbf{y}\|_K = \|\mathbf{x}\|_K + \|\mathbf{y}\|_K < 1 + 1 = 2,$$

but this is a contradiction since by assumption $\lambda_1(K, L) \geq 2$. To see that T is measure preserving, we note that T is injective and acts on K by a finite number of translations (each of which preserve Lebesgue measure).

To build the map for general $t > 0$, we define the map $T_t : tK' \rightarrow t\mathcal{V}$ by $T_t(\mathbf{x}) = tT(\mathbf{x}/t)$. It is simple to check that the generalized map T_t satisfies all the required properties.

Claim 2: $\forall t > 0, \gamma_s(tK) \leq \gamma_s(t\mathcal{V})$.

We first note that $\gamma_s(tK) = \gamma_s(tK')$. From the previous claim, we have the existence of the map $T : tK' \rightarrow t\mathcal{V}$ with the aforementioned properties. From here we see that

$$\begin{aligned} \frac{1}{s^n} \int_{t\mathcal{V}} e^{-\pi\|\mathbf{x}/s\|_2^2} d\mathbf{x} &\geq \frac{1}{s^n} \int_{T(tK')} e^{-\pi\|\mathbf{x}/s\|_2^2} d\mathbf{x} \\ &= \frac{1}{s^n} \int_{tK'} e^{-\pi\|T(\mathbf{x})/s\|_2^2} d\mathbf{x} \geq \frac{1}{s^n} \int_{tK'} e^{-\pi\|\mathbf{x}/s\|_2^2} d\mathbf{x} = \gamma_s(tK') \end{aligned}$$

as needed.

The lemma thus follows. □

Corollary 6.3.7. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a lattice, and let $K \subseteq \mathbb{R}^n$ be a symmetric convex body. Then for $p > 0$, we have that*

$$\frac{\mathbb{E}[\|X\|_{\mathcal{V}}^p]^{\frac{1}{p}}}{2} \leq \frac{\mathbb{E}[\|X\|_K^p]^{\frac{1}{p}}}{\lambda_1(K, \mathcal{L})}$$

where X is distributed as $D_{s, \mathbf{0}}$, $s > 0$, and $\mathcal{V} = \mathcal{V}(\mathcal{L})$.

Proof. For $x \in \mathbb{R}^n$ and $t > 0$, note that $\|x\|_{tK} = \frac{1}{t}\|x\|_K$. Hence $\lambda_1(tK, \mathcal{L}) = \frac{1}{t}\lambda_1(K, \mathcal{L})$ and $\mathbb{E}[\|X\|_{tK}^p]^{\frac{1}{p}} = \frac{1}{t}\mathbb{E}[\|X\|_K^p]^{\frac{1}{p}}$. From this, we get that the above inequality is invariant under scalings of K . It therefore suffices to prove the inequality when $\lambda_1(K, \mathcal{L}) = 2$, which follows immediately from Lemma 6.3.6 (setting $f(z) = z^p$ for $z \in \mathbb{R}_+$). □

We are now ready to give a the proof of the main smoothing parameter bound.

Proof of Theorem 6.3.1. Let Y be distributed as $D_{\mathcal{L}^*, s, \mathbf{0}}$ and $\mathcal{V} = \mathcal{V}(\mathcal{L}^*)$. Noting that

$$\Pr[Y \neq \mathbf{0}] = \frac{\rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\})}{1 + \rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\})} \text{ we get that } \rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \epsilon \Leftrightarrow \Pr[Y \neq \mathbf{0}] \leq \frac{\epsilon}{1 + \epsilon}$$

Now using Lemma 6.3.5 and Markov's inequality, we have that

$$\Pr[Y \neq \mathbf{0}] = \frac{\rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\})}{\rho_{\frac{1}{s}}(\mathcal{L}^*)} \leq 1 - \gamma_{\frac{1}{s}}(\mathcal{V}) = \Pr\left[\frac{1}{s}X \in \mathbb{R}^n \setminus \mathcal{V}\right] = \Pr\left[\left\|\frac{1}{s}X\right\|_{\mathcal{V}} \geq 1\right] \leq \mathbb{E}\left[\left\|\frac{1}{s}X\right\|_{\mathcal{V}}^p\right]$$

From Corollary 6.3.7 we have that $E[\|\frac{1}{s}X\|_{\mathcal{V}}^p] \leq 2^p \frac{E[\|\frac{1}{s}X\|_K^p]}{\lambda_1(K, \mathcal{L}^*)}$. Hence, combining the above inequalities together we get that

$$\Pr[Y \neq \mathbf{0}] \leq E[\|\frac{1}{s}X\|_{\mathcal{V}}^p] \leq 2^p \frac{E[\|\frac{1}{s}X\|_K^p]}{\lambda_1(K, \mathcal{L}^*)^p} = \left(\frac{2}{s}\right)^p \frac{E[\|X\|_K^p]}{\lambda_1(K, \mathcal{L}^*)^p}$$

Therefore to insure $\rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \epsilon$, it suffices to choose $s > 0$ satisfying

$$\frac{\epsilon}{1 + \epsilon} \leq \left(\frac{2}{s}\right)^p \frac{E[\|X\|_K^p]}{\lambda_1(K, \mathcal{L}^*)^p} \Leftrightarrow s \geq 2 \left(1 + \frac{1}{\epsilon}\right)^{\frac{1}{p}} \frac{E[\|X\|_K^p]^{\frac{1}{p}}}{\lambda_1(K, \mathcal{L}^*)}$$

The desired bound on $\eta_{\epsilon}(\mathcal{L})$ thus follows. \square

To prove Theorem 6.3.2, we will use the following theorem of Latala and Oleszkiewicz [81].

While a far weaker and elementary bound would suffice, the following theorem will allow us to derive better constants.

Theorem 6.3.8. *Let $X \in \mathbb{R}^n$ be an n dimensional Gaussian random variable. Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body, and let $\alpha \geq 0$ be chosen such that $\Pr[X \in K] = \Pr[|X_1| \leq \alpha]$. Then the following holds:*

- For $t \in [0, 1]$, $\Pr[X \in tK] \leq \Pr[|X_1| \leq t\alpha]$.
- For $t \geq 1$, $\Pr[X \in tK] \geq \Pr[|X_1| \leq t\alpha]$.

Proof of Theorem 6.3.2. The upper bound follows directly from Theorem 6.3.1, so we need only prove the lower bound. Define $m > 0$ to satisfy $\Pr[X \in mV] = \Pr[\|X\|_{\mathcal{V}} \leq m] = \frac{1}{2}$. Define α by the relation $\Pr[|X_1| \leq \alpha] = \frac{1}{2}$. Then by Theorem 6.3.8, we have that $\Pr[\|X\|_{\mathcal{V}} \geq tm] \leq \Pr[|X_1| \geq t\alpha]$ for $t \geq 1$. We now have that

$$\begin{aligned} E[\|X\|_{\mathcal{V}}] &= \int_0^{\infty} \Pr[\|X\|_{\mathcal{V}} \geq t] dt \leq m + \int_m^{\infty} \Pr[\|X\|_{\mathcal{V}} \geq t] dt \leq m + \int_m^{\infty} \Pr[|X_1| \geq \frac{\alpha t}{m}] dt \\ &= \left(1 + \frac{1}{\alpha} \int_{\alpha}^{\infty} \Pr[|X_1| \geq t] dt\right)m = \left(1 + \frac{2}{\alpha} \int_0^{\infty} te^{-\pi(t+\alpha)^2} dt\right)m = (1 + \beta)m \end{aligned}$$

Let $s = \eta_{\frac{1}{2}}(\mathcal{L})$. Since $\rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\}) = \frac{1}{2}$, by Lemma 6.3.5, we have that

$$1 - \gamma_{\frac{1}{2s}}(\mathcal{V}) \leq \rho_{\frac{1}{s}}(\mathcal{L}^* \setminus \{\mathbf{0}\}) = \frac{1}{2}$$

Therefore $\gamma_{\frac{1}{2s}}(\mathcal{V}) = \Pr[X \in 2s\mathcal{V}] \geq \frac{1}{2}$. Since $\Pr[X \in mV] = \frac{1}{2}$, we have that $2s \geq m \geq \frac{1}{1+\beta} \mathbb{E}[\|X\|_V]$. Hence

$$\eta_{\frac{1}{2}}(\mathcal{L}) = s \geq \frac{1}{2(1+\beta)} \mathbb{E}[\|X\|_V].$$

Using numerical approximations one gets that $\beta \leq .444$, and hence $\frac{1}{2(1+\beta)} \geq \frac{1}{3}$ as needed. \square

6.4 Comparing the Discrete and Continuous Gaussian

As mentioned in the preliminaries, the smoothing parameter measures how close the discrete Gaussian is to the continuous Gaussian. In this section, we first relate known results which show that the convolution of a “smooth” discrete Gaussian with a continuous Gaussian is close to Gaussian, and the a convolution of discrete Gaussians is again close to discrete Gaussian. We shall then use these results to give nearly matching upper and lower bounds on the expected norm of the discrete Gaussian under general norms.

The first fundamental lemma of this section, which first appeared in [90], gives tight bounds on discrete Gaussian sums above the smoothing parameter.

Lemma 6.4.1. *Let \mathcal{L} be an n dimensional lattice, and let $s \geq \eta_\epsilon(\mathcal{L})$ for $0 < \epsilon < 1$.*

Then $\forall \mathbf{x} \in \mathbb{R}^n$,

$$(1 - \epsilon) \frac{s^n}{\det(\mathcal{L})} \leq \rho_s(\mathcal{L} + \mathbf{x}) \leq (1 + \epsilon) \frac{s^n}{\det(\mathcal{L})}$$

Proof. By the Poisson summation formula

$$\rho_s(\mathcal{L} + \mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{L}} \rho_s(\mathbf{y} + \mathbf{x}) = \frac{s^n}{\det(\mathcal{L})} \sum_{\mathbf{y} \in \mathcal{L}^*} e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} \rho_{\frac{1}{s}}(\mathbf{y})$$

Since $s \geq \eta_\epsilon(\mathcal{L})$, we have that

$$\begin{aligned} \left| \left(\frac{s^n}{\det(\mathcal{L})} \sum_{\mathbf{y} \in \mathcal{L}^*} e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} \rho_{\frac{1}{s}}(\mathbf{y}) \right) - \frac{s^n}{\det(\mathcal{L})} \right| &= \frac{s^n}{\det(\mathcal{L})} \left| \sum_{\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} e^{2\pi i \langle \mathbf{x}, \mathbf{y} \rangle} \rho_{\frac{1}{s}}(\mathbf{y}) \right| \\ &\leq \frac{s^n}{\det(\mathcal{L})} \left| \sum_{\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \rho_{\frac{1}{s}}(\mathbf{y}) \right| \leq \epsilon \frac{s^n}{\det(\mathcal{L})} \end{aligned}$$

as needed. \square

We first relate theorems about convolutions of discrete Gaussians and continuous Gaussians. The following is a combination of Claim 3.9 from [108] and Theorem 3.1 from [102].

Theorem 6.4.2. *Let $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^n$, and $s_1, s_2 > 0$ such that $\frac{s_1 s_2}{\sqrt{s_1^2 + s_2^2}} \geq \eta_\epsilon(\mathcal{L})$, $\epsilon > 0$. Setting $\bar{\mathbf{a}} = \mathbf{a}_1 + \mathbf{a}_2$, and $\bar{s} = \sqrt{s_1^2 + s_2^2}$, then for the random variables $X \sim D_{\mathcal{L} + \mathbf{a}_1, s_1}$ and $Y \sim D_{\mathcal{L} + \mathbf{a}_2, s_2}$ we have that*

$$\left(\frac{1-\epsilon}{1+\epsilon}\right)^2 D_{\mathcal{L} + \bar{\mathbf{a}}, \bar{s}}(\mathbf{x}) \leq \Pr[X + Y = \mathbf{x}] \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)^2 D_{\mathcal{L} + \bar{\mathbf{a}}, \bar{s}}(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^n$$

Furthermore, if $Y \sim D_{s_2}$, we have that

$$\left(\frac{1-\epsilon}{1+\epsilon}\right) \frac{\rho_{\bar{s}}(x)}{\bar{s}^n} \leq \text{dPr}[X + Y = \mathbf{x}] \leq \left(\frac{1+\epsilon}{1-\epsilon}\right) \frac{\rho_{\bar{s}}(\mathbf{x})}{\bar{s}^n} \quad \forall \mathbf{x} \in \mathbb{R}^n$$

Proof. We prove the first claim. Since X, Y are supported on $\mathcal{L} + \mathbf{a}_1, \mathcal{L} + \mathbf{a}_2$ respectively, we see that $X + Y \in \mathcal{L} + \mathbf{a}_1 + \mathbf{a}_2 = \mathcal{L} + \bar{\mathbf{a}}$. For $\mathbf{x} \in \mathcal{L} + \bar{\mathbf{a}}$, we have that

$$\Pr[X + Y = \mathbf{x}] = \frac{1}{\rho_{s_1}(\mathcal{L} + \mathbf{a}_1)} \frac{1}{\rho_{s_2}(\mathcal{L} + \mathbf{a}_2)} \sum_{\mathbf{y} \in \mathcal{L} + \mathbf{a}_2} \rho_{s_1}(\mathbf{x} - \mathbf{y}) \rho_{s_2}(\mathbf{y})$$

Letting $\tilde{s} = \frac{s_1 s_2}{\sqrt{s_1^2 + s_2^2}}$, we have that

$$\begin{aligned} \sum_{\mathbf{y} \in \mathcal{L} + \mathbf{a}_2} \rho_{s_1}(\mathbf{x} - \mathbf{y}) \rho_{s_2}(\mathbf{y}) &= \sum_{\mathbf{y} \in \mathcal{L} + \mathbf{a}_2} \exp \left[-\pi \left(\left\| \frac{\mathbf{x} - \mathbf{y}}{s_1} \right\|_2^2 + \left\| \frac{\mathbf{y}}{s_2} \right\|_2^2 \right) \right] \quad (\text{Completing the square}) \\ &= \sum_{\mathbf{y} \in \mathcal{L} + \mathbf{a}_2} \exp \left[-\pi \left(\left\| \frac{\mathbf{x}}{\sqrt{s_1^2 + s_2^2}} \right\|_2^2 + \left\| \frac{\sqrt{s_1^2 + s_2^2}}{s_1 s_2} \left(\mathbf{y} - \frac{s_2^2}{s_1^2 + s_2^2} \mathbf{x} \right) \right\|_2^2 \right) \right] \\ &= \rho_{\bar{s}}(\mathbf{x}) \rho_{\bar{s}} \left(\mathcal{L} + \mathbf{a}_2 - \frac{s_2^2}{s_1^2 + s_2^2} \mathbf{x} \right) \end{aligned}$$

Since $s_1, s_2, \tilde{s} \geq \eta_\epsilon(\mathcal{L})$, by Lemma 6.4.1 we have that

$$\begin{aligned} \frac{\rho_{\bar{s}} \left(\mathcal{L} + \mathbf{a}_2 - \frac{s_2^2}{s_1^2 + s_2^2} \mathbf{x} \right)}{\rho_{s_1}(\mathcal{L} + \mathbf{a}_1) \rho_{s_2}(\mathcal{L} + \mathbf{a}_2)} &\leq \frac{(1+\epsilon)}{(1-\epsilon)^2} \frac{\tilde{s}^n}{s_1^n s_2^n} \det(\mathcal{L}) \\ &= \frac{(1+\epsilon)}{(1-\epsilon)^2} \frac{\det(\mathcal{L})}{\bar{s}^n} \leq \frac{(1+\epsilon)^2}{(1-\epsilon)^2} \frac{1}{\rho_{\bar{s}}(\mathcal{L} + \bar{\mathbf{a}})} \end{aligned}$$

For the lower bound, we similarly derive that that

$$\frac{\rho_{\bar{s}}\left(\mathcal{L} + \mathbf{a}_2 - \frac{s_2^2}{s_1^2 + s_2^2} \mathbf{x}\right)}{\rho_{s_1}(\mathcal{L} + \mathbf{a}_1)\rho_{s_2}(\mathcal{L} + \mathbf{a}_2)} \geq \frac{(1 - \epsilon)^2}{(1 + \epsilon)^2} \frac{1}{\rho_{\bar{s}}(\mathcal{L} + \bar{\mathbf{a}})}$$

Combining all the above inequalities, we get that

$$\frac{(1 - \epsilon)^2}{(1 + \epsilon)^2} \frac{\rho_{\bar{s}}(\mathbf{x})}{\rho_{\bar{s}}(\mathcal{L} + \bar{\mathbf{a}})} \leq \Pr[X + Y = \mathbf{x}] \leq \frac{(1 + \epsilon)^2}{(1 - \epsilon)^2} \frac{\rho_{\bar{s}}(\mathbf{x})}{\rho_{\bar{s}}(\mathcal{L} + \bar{\mathbf{a}})}$$

Now let Y be distributed as D_{s_2} . Then the density of $X + Y$, for $\mathbf{x} \in \mathbb{R}^n$, satisfies

$$\mathrm{dPr}[X + Y = \mathbf{x}] = \frac{1}{\rho_{s_1}(\mathcal{L} + \mathbf{a}_1)s_2^n} \sum_{\mathbf{y} \in \mathcal{L} + \mathbf{a}_1} \rho_{s_1}(\mathbf{y})\rho_{s_2}(\mathbf{x} - \mathbf{y})$$

Then using the identical analysis as above, we get that

$$\sum_{\mathbf{y} \in \mathcal{L} + \mathbf{a}_1} \rho_{s_1}(\mathbf{y})\rho_{s_2}(\mathbf{x} - \mathbf{y}) = \rho_{\bar{s}}(\mathbf{x})\rho_{\bar{s}}\left(\mathcal{L} + \mathbf{a}_1 - \frac{s_2^2}{s_1^2 + s_2^2} \mathbf{x}\right).$$

In the same way as in the previous analysis, we get that

$$\frac{(1 - \epsilon)}{(1 + \epsilon)} \frac{1}{\bar{s}^n} \leq \frac{\rho_{\bar{s}}\left(\mathcal{L} + \mathbf{a}_1 - \frac{s_2^2}{s_1^2 + s_2^2} \mathbf{x}\right)}{\rho_{s_1}(\mathcal{L} + \mathbf{a}_1)s_2^n} \leq \frac{(1 + \epsilon)}{(1 - \epsilon)} \frac{1}{\bar{s}^n}$$

Combining the above inequalities yields the result. \square

The next theorem shows nearly tight bounds for discrete Gaussian norm expectations beyond the smoothing parameter. Compared to previous literature, our main contribution is that we give nearly matching upper and lower bounds for discrete Gaussian expectations under general norms.

Theorem 6.4.3. *Let $K \subseteq \mathbb{R}^n$ be a convex body containing the origin in its interior and let \mathcal{L} be an n dimensional lattice. For $\delta > 1$, $s \geq \delta\eta_\epsilon(\mathcal{L})$, for any $\mathbf{c} \in \mathbb{R}^n$ and $p \geq 1$ the random variables $Y \sim D_{\mathcal{L} + \mathbf{c}, s}$ and $X \sim D_{s, \mathbf{0}}$ satisfy*

$$\mathbb{E}[\|Y\|_K^p]^{\frac{1}{p}} \leq \left(\frac{1 + \epsilon}{1 - \epsilon}\right)^{\frac{1}{p}} \sqrt{\frac{\delta^2}{\delta^2 - 1}} \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}}$$

Furthermore, as long as $\delta \geq 2\left(\frac{1 + \epsilon}{1 - \epsilon}\right)^{\frac{3}{p}}$, we have that

$$\mathbb{E}[\|Y\|_K^p]^{\frac{1}{p}} \geq \left(\frac{1 - \epsilon}{1 + \epsilon}\right)^{\frac{2}{p}} \left(1 - \left(\frac{1 + \epsilon}{1 - \epsilon}\right)^{\frac{3}{p}} \frac{2}{\delta}\right) \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}}$$

Proof. We begin by proving the upper bound. Let Z be distributed as $D_{\bar{s}, \mathbf{0}}$, for $\bar{s} = \sqrt{\frac{\delta^2}{\delta^2-1}}\eta_\epsilon(\mathcal{L})$. Since the gauge $\|\cdot\|_K$ is convex and non-negative, and $p \geq 1$, by Jensen's inequality we have that

$$\begin{aligned} \mathbb{E}[\|Y + Z\|_K^p] &= \sum_{\mathbf{y} \in \mathcal{L}} \mathbb{E}[\|\mathbf{y} + Z\|_K^p] \Pr(Y = \mathbf{y}) \geq \sum_{\mathbf{y} \in \mathcal{L}} \|\mathbf{y} + \mathbb{E}[Z]\|_K^p \Pr(Y = \mathbf{y}) \\ &= \sum_{\mathbf{y} \in \mathcal{L}} \|\mathbf{y}\|_K^p \Pr(Y = \mathbf{y}) = \mathbb{E}[\|Y\|_K^p] \end{aligned}$$

Now note that

$$\frac{\bar{s}s}{\sqrt{\bar{s}^2 + s^2}} = \frac{1}{\sqrt{\frac{1}{\bar{s}^2} + \frac{1}{s^2}}} = \frac{1}{\sqrt{\frac{\delta^2-1}{\delta^2} + \frac{1}{\delta^2}}}\eta_\epsilon(\mathcal{L}) = \eta_\epsilon(\mathcal{L})$$

and that $\sqrt{\bar{s}^2 + s^2} = \sqrt{\frac{\delta^4}{\delta^2-1}}\eta_\epsilon(\mathcal{L}) = s\sqrt{\frac{\delta^2}{\delta^2-1}}$. Letting $\bar{\delta} = \sqrt{\frac{\delta^2}{\delta^2-1}}$ by proposition 6.4.2 we have that

$$\frac{1-\epsilon}{1+\epsilon} \frac{\rho_{\bar{\delta}s}(\mathbf{x})}{(\bar{\delta}s)^n} \leq \text{dPr}[Y + Z = \mathbf{x}] \leq \frac{1+\epsilon}{1-\epsilon} \frac{\rho_{\bar{\delta}s}(\mathbf{x})}{(\bar{\delta}s)^n} \quad \forall \mathbf{x} \in \mathbb{R}^n$$

Therefore

$$\begin{aligned} \mathbb{E}[\|Y\|_K^p]^{\frac{1}{p}} &\leq \mathbb{E}[\|Y + Z\|_K^p]^{\frac{1}{p}} = \left(\int_{\mathbb{R}^n} \|\mathbf{x}\|_K^p \text{dPr}[Y + Z = \mathbf{x}] \right)^{\frac{1}{p}} \\ &\leq \left(\int_{\mathbb{R}^n} \|\mathbf{x}\|_K \frac{1+\epsilon}{1-\epsilon} \frac{\rho_{\bar{\delta}s}(\mathbf{x})}{(\bar{\delta}s)^n} \text{d}\mathbf{x} \right)^{\frac{1}{p}} = \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} \bar{\delta} \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}} \end{aligned}$$

as needed.

We now prove the lower bound. Redefine \bar{s} to be $\sqrt{\frac{\delta^2 + (c^2-1)\delta^4}{\delta^2-1}}\eta_\epsilon(\mathcal{L})$ (c will be chosen later), $c \geq 1$, where we note that now $\sqrt{s^2 + \bar{s}^2} = c\sqrt{\frac{\delta^2}{\delta^2-1}}s = c\bar{\delta}s$. Let $t = c\bar{\delta}s$.

Let F denote a fundamental parallelepiped of \mathcal{L} (i.e. $F = B[0, 1]^n$ for a basis

matrix B of \mathcal{L}). Then

$$\begin{aligned}
\mathbb{E}[\|c\bar{\delta}X\|_K^p] &= \frac{1}{t^n} \int_{\mathbb{R}^n} \|\mathbf{x}\|_K^p \rho_t(\mathbf{x}) d\mathbf{x} = \frac{1}{t^n} \int_F \sum_{\mathbf{y} \in \mathcal{L} + \mathbf{x}} \|\mathbf{y}\|_K^p \rho_t(\mathbf{y}) d\mathbf{x} \\
&= \frac{1}{t^n} \int_F \rho_t(\mathcal{L} + \mathbf{x}) \left(\sum_{\mathbf{y} \in \mathcal{L} + \mathbf{x}} \|\mathbf{y}\|_K^p \frac{\rho_t(\mathbf{y})}{\rho_t(\mathcal{L} + \mathbf{x})} \right) d\mathbf{x} \\
&\leq \max_{\mathbf{x} \in F} \left(\sum_{\mathbf{y} \in \mathcal{L} + \mathbf{x}} \|\mathbf{y}\|_K^p \frac{\rho_t(\mathbf{y})}{\rho_t(\mathcal{L} + \mathbf{x})} \right)
\end{aligned} \tag{6.4.1}$$

where the inequality follows since $t^{-n} \rho_t(\mathcal{L} + \mathbf{x})$ induces a probability distribution over F .

Let \mathbf{x}^* denote an element of F maximizing the expression in Equation (6.4.1). Let $W \sim D_{\mathcal{L} + \mathbf{x}^*, t, \mathbf{0}}$. Then by equation (6.4.1) and the choice of \mathbf{x}^* , we have that $\mathbb{E}[\|W\|_K^p] \geq \mathbb{E}[\|c\bar{\delta}X\|_K^p]$. Now let us redefine Z to be distributed as $D_{\mathcal{L} + \mathbf{x}^* - \mathbf{c}, \bar{s}, \mathbf{0}}$. Since $\frac{\bar{s}s}{\sqrt{\bar{s}^2 + s^2}} \geq \eta_\epsilon(\mathcal{L})$ and $\mathbf{c} + \mathbf{x}^* - \mathbf{c} = \mathbf{x}^*$, by Proposition 6.4.2 we have that

$$\left(\frac{1 - \epsilon}{1 + \epsilon} \right)^2 D_{\mathcal{L} + \mathbf{x}^*, t, \mathbf{0}}(\mathbf{x}) \leq \Pr[Y + Z = \mathbf{x}] \leq \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^2 D_{\mathcal{L} + \mathbf{x}^*, t, \mathbf{0}}(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^n$$

Therefore we have that

$$\begin{aligned}
\mathbb{E}[\|Y + Z\|_K^p]^{\frac{1}{p}} &= \left(\sum_{\mathbf{w} \in \mathcal{L} + \mathbf{x}^*} \|\mathbf{w}\|_K^p \Pr[Y + Z = \mathbf{w}] \right)^{\frac{1}{p}} \geq \left(\sum_{\mathbf{w} \in \mathcal{L} + \mathbf{x}^*} \|\mathbf{w}\|_K^p \left(\frac{1 - \epsilon}{1 + \epsilon} \right)^2 \frac{\rho_t(\mathbf{w})}{\rho_t(\mathcal{L} + \mathbf{x}^*)} \right)^{\frac{1}{p}} \\
&= \left(\frac{1 - \epsilon}{1 + \epsilon} \right)^{\frac{2}{p}} \mathbb{E}[\|W\|_K^p]^{\frac{2}{p}} \geq \left(\frac{1 - \epsilon}{1 + \epsilon} \right)^{\frac{2}{p}} \mathbb{E}[\|c\bar{\delta}X\|_K^p]^{\frac{1}{p}}
\end{aligned} \tag{6.4.2}$$

Using Holder's inequality (and the fact that $\|\cdot\|_K$ satisfies the triangle inequality), we have that

$$\mathbb{E}[\|Y\|_K^p]^{\frac{1}{p}} + \mathbb{E}[\|Z\|_K^p]^{\frac{1}{p}} \geq \mathbb{E}[\|Y + Z\|_K^p]^{\frac{1}{p}} \tag{6.4.3}$$

Using the first part of the Theorem, a straightforward computations reveals that

$$\mathbb{E}[\|Z\|_K^p]^{\frac{1}{p}} \leq \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^{\frac{1}{p}} \frac{1 + (c^2 - 1)\delta^2}{\sqrt{1 + (c^2 - 1)\delta^4}} \bar{\delta} \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}} \tag{6.4.4}$$

Combining Equation (6.4.2), (6.4.3), (6.4.4), we have that

$$\mathbb{E}[\|Y\|_K^p]^{\frac{1}{p}} \geq \left(\left(\frac{1 - \epsilon}{1 + \epsilon} \right)^{\frac{2}{p}} c - \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^{\frac{1}{p}} \frac{1 + (c^2 - 1)\delta^2}{\sqrt{1 + (c^2 - 1)\delta^4}} \right) \bar{\delta} \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}}$$

Setting $c = \sqrt{1 + \frac{1}{\delta^2}}$, we get that

$$\begin{aligned} \mathbb{E}[\|Y\|_K^p]^{\frac{1}{p}} &\geq \left(\left(\frac{1-\epsilon}{1+\epsilon} \right)^{\frac{2}{p}} \sqrt{1 + \frac{1}{\delta^2}} - \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} \frac{2}{\sqrt{1 + \delta^2}} \right) \bar{\delta} \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}} \\ &\geq \left(\left(\frac{1-\epsilon}{1+\epsilon} \right)^{\frac{2}{p}} - \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} \frac{2}{\delta} \right) \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}} \\ &\geq \left(\frac{1-\epsilon}{1+\epsilon} \right)^{\frac{2}{p}} \left(1 - \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{3}{p}} \frac{2}{\delta} \right) \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}} \end{aligned}$$

as long as $\delta \geq 2 \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{3}{p}}$. □

6.5 Flatness Theorem Proof

We are now ready to prove the flatness theorem. We give two versions of the theorem, one with respect to the smoothing parameter, and one with respect to the ℓ^* estimate.

Theorem 6.5.1. *Let $K \subseteq \mathbb{R}^n$ be a convex body and let $\mathcal{L} \subseteq \mathbb{R}^n$ be an n dimensional lattice. Then for all $\epsilon \in (0, 1)$ and $p \geq 1$, we have that*

$$\mu(K, L) \leq \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} \sqrt{\frac{2}{\pi}} \eta_\epsilon(\mathcal{L}) \inf_{\mathbf{x} \in K} \ell_{K-\mathbf{x}}^p(I_n) \quad (6.5.1)$$

and that

$$\mu(K, \mathcal{L}) \lambda_1((K - K)^*, \mathcal{L}^*) \leq \frac{4}{\pi} 8^{\frac{1}{p}} \inf_{\substack{T \in \mathbb{R}^{n \times n}, \det(T)=1 \\ \mathbf{x} \in K}} \ell_{K-\mathbf{x}}^p(T) \ell_{(K-\mathbf{x})^*}^p(T) \quad (6.5.2)$$

Proof. We begin by proving inequality (6.5.1). To bound $\mu(K, \mathcal{L})$ we must find $t \geq 0$ such that $\forall \mathbf{c} \in \mathbb{R}^n, \mathbf{c} + tK \cap \mathcal{L} \neq \emptyset$. Since the covering radius is shift invariant, we may center K arbitrarily (i.e. shift K to $K - \mathbf{x}$ for some $\mathbf{x} \in K$). Hence we may assume that the origin in is in the interior of K .

Pick $\mathbf{c} \in \mathbb{R}^n$. Let $s = \sqrt{2}\eta_\epsilon(\mathcal{L})$. Let $Y_{\mathbf{c}}$ be distributed as $D_{\mathcal{L},s,\mathbf{c}}$, and let X be distributed as $D_{1,\mathbf{0}}$. Note that $Y_{\mathbf{c}} - \mathbf{c} \sim D_{\mathcal{L}-\mathbf{c},s,\mathbf{0}}$, therefore by Theorem 6.4.3

$$\begin{aligned} \mathbb{E}[\|Y_{\mathbf{c}} - \mathbf{c}\|_K^p]^{\frac{1}{p}} &\leq \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} \sqrt{2} \mathbb{E}[\|sX\|_K^p]^{\frac{1}{p}} = \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} 2\eta_\epsilon(\mathcal{L}) \mathbb{E}[\|X\|_K^p]^{\frac{1}{p}} \\ &= \left(\frac{1+\epsilon}{1-\epsilon} \right)^{\frac{1}{p}} \sqrt{\frac{2}{\pi}} \eta_\epsilon(\mathcal{L}) \ell_K^p(I_n) \end{aligned}$$

Letting $t = \mathbb{E}[\|Y_{\mathbf{c}} - \mathbf{c}\|_K^p]^{\frac{1}{p}}$, we note that there must a lattice point $\mathbf{y} \in \mathcal{L}$ such that $\|\mathbf{x} - \mathbf{c}\|_K \leq t \Rightarrow \mathbf{c} + tK \cap \mathcal{L} \neq \emptyset$. Therefore

$$\mu(K, \mathcal{L}) \leq \max_{\mathbf{c} \in \mathbb{R}^n} \mathbb{E}[\|Y_{\mathbf{c}} - \mathbf{c}\|_K] \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)^{\frac{1}{p}} \sqrt{\frac{2}{\pi}} \eta_{\epsilon}(\mathcal{L}) \ell_K^p(I_n) \quad (6.5.3)$$

Lastly, since we have freedom in where we center K , we may replace $\ell_K^p(I_n)$ above by $\inf_{\mathbf{x} \in K} \ell_{K-\mathbf{x}}^p(I_n)$.

We now derive inequality 6.5.2 from inequality 6.5.1. Since $(K - K)^*$ is a symmetric convex body by Theorem 6.3.1, we have that

$$\eta_{\epsilon}(\mathcal{L}) \leq 2 \left(1 + \frac{1}{\epsilon}\right)^{\frac{1}{p}} \frac{\mathbb{E}[\|X\|_{(K-K)^*}^p]^{\frac{1}{p}}}{\lambda_1((K - K)^*, \mathcal{L}^*)} \quad (6.5.4)$$

Since we may assume that $\mathbf{0} \in K$, we have that for $\mathbf{x} \in \mathbb{R}^n$ that $\|\mathbf{x}\|_{(K-K)^*} = \|\mathbf{x}\|_{K^*} + \|\mathbf{x}\|_{K^*}$. Therefore by Minkowski's inequality and the symmetry of the Gaussian we have that

$$\begin{aligned} \mathbb{E}[\|X\|_{(K-K)^*}^p]^{\frac{1}{p}} &= \mathbb{E}[(\|X\|_{K^*} + \|\mathbf{x}\|_{K^*})^p]^{\frac{1}{p}} \leq \mathbb{E}[\|X\|_{K^*}^p]^{\frac{1}{p}} + \mathbb{E}[\|\mathbf{x}\|_{K^*}^p]^{\frac{1}{p}} \\ &= 2 \mathbb{E}[\|X\|_{K^*}^p]^{\frac{1}{p}} = \sqrt{\frac{2}{\pi}} \ell_{K^*}^p(I_n) \end{aligned} \quad (6.5.5)$$

Now combining Equations (6.5.4), (6.5.5) and (6.5.3), we get that

$$\mu(K, \mathcal{L}) \lambda_1((K - K)^*, \mathcal{L}^*) \leq \frac{4}{\pi} \left(\frac{(1+\epsilon)^2}{(1-\epsilon)\epsilon}\right)^{\frac{1}{p}} \ell_K^p(I_n) \ell_{K^*}^p(I_n)$$

Minimizing over ϵ , we get that $\left(\frac{(1+\epsilon)^2}{(1-\epsilon)\epsilon}\right)$ attains a minimum value of 8 at $\epsilon = \frac{1}{3}$. Now we note that the left hand side is invariant under affine transformations, i.e. for $T \in \mathbb{R}^{n \times n}$, $\det(T) = 1$ and $\mathbf{x} \in K$, we have that

$$\begin{aligned} \mu(T(K - \mathbf{x}), T\mathcal{L}) \lambda_1((T(K - \mathbf{x}) - T(K - \mathbf{x}))^*, (T\mathcal{L})^*) &= \mu(TK, T\mathcal{L}) \lambda_1((T(K - K))^*, (T\mathcal{L})^*) \\ &= \mu(K, L) \lambda_1((K - K)^*, \mathcal{L}^*) \end{aligned}$$

Therefore, we may minimize over all such choices over the right hand side, which yields

$$\mu(K, L) \lambda_1((K - K)^*, \mathcal{L}^*) \leq \frac{4}{\pi} 8^{\frac{1}{p}} \inf_{\substack{T \in \mathbb{R}^{n \times n}, \det(T)=1 \\ \mathbf{x} \in K}} \ell_{K-\mathbf{x}}^p(T) \ell_{(K-\mathbf{x})^*}^p(T)$$

as claimed. □

6.6 Conclusion

Kinchine’s flatness theorem is a fundamental result in the geometry of numbers. From the algorithmic perspective, it is the key structural result on lattices responsible for the efficiency of the current fastest algorithms for Integer Programming (see Chapter 7 for details). From the complexity perspective, it has been a valuable tool for relating the approximation complexity of various lattice problems. Within Cryptography, perhaps more important than the theorem itself has been the tools developed to prove it, i.e. the discrete Gaussian distribution and its associated properties [10, 11, 12]. The properties of the discrete Gaussian have played a central role in the development of modern lattice based cryptography, and are currently behind the tightest worst-case to average case reductions for lattice problems [90, 108].

In this chapter, we have come full circle, by bringing to bear the insights developed within the cryptographic study of the discrete Gaussian back onto the original problem it was designed to solve, i.e. obtaining near-optimal bounds for the flatness theorem. For our main result, we have improved on a theorem of Banaszczyk by giving a tighter reduction (with very small constants) from bounding the flatness constant to bounding the ℓ^* estimate in convex geometry. Our main improvement here was to avoid Banaszczyk’s use of Talagrand’s majorizing measure theorem, by giving a new geometric characterization of the smoothing parameter in terms of the Gaussian measure of the voronoi cell. We note that this characterization is crucially used in [27] to show that the problem of estimating the smoothing parameter within a factor $2 + o(1)$ is in SZK (Statistical Zero Knowledge). Our second contribution, was to show that above the smoothing parameter, the expected norm of the discrete Gaussian is both upper and lower bounded by the corresponding Gaussian expectation. The main novelty here is the lower bound for general norms (though it is not used in the proof of the flatness theorem), which had only previously been shown for the ℓ_2 norm [90].

Future Research. As mentioned in the introduction, the current bounds on the ℓ^* estimate for asymmetric convex bodies, i.e. $O(n^{\frac{4}{3}} \text{polylog } n)$, are very far from the current lower bound of $O(n \log n)$. Closing this gap remains a major open problem in convex geometry.

A second line of research, involves understanding finer properties of the smoothing parameter. We note that the presented proof of the flatness theorem simply utilizes the shortest vector of the dual lattice as a proxy for the smoothing parameter. Hence if we are interested in connecting the smoothing parameter directly to other lattice quantities, it maybe possible to get tighter connection factors. Furthermore, the current proofs only use discrete Gaussians with spherical covariances. A natural question is whether one might be able to prove stronger results by directly optimizing over the discrete Gaussian covariances.

Lastly, another natural question is whether one can prove stronger transference type theorems by using non-Gaussian densities functions.

CHAPTER VII

THE INTEGER PROGRAMMING PROBLEM

The Integer Programming Problem (IP), i.e. the problem of deciding whether a convex set contains an integer point, is a fundamental problem in Computer Science and Operations Research. In this chapter, we present two improved algorithms for IP feasibility, the first being based on a framework developed by Lenstra [84], and second on a framework developed by Kannan [70]. For an n -variable integer program, our Lenstra type and Kannan type algorithm respectively run in $2^{O(n)}(n^{\frac{4}{3}} \text{polylog}(n))^n$ and $2^{O(n)}n^n$ time, and both use 2^n space. Together they represent the fastest algorithms of either type, and the fastest known algorithms for IP feasibility in general. Lastly, we generalize these algorithms to their natural optimization counterpart, by giving a randomized $2^{O(n)}n^n$ expected time and 2^n space algorithm to minimize a general convex function over the integer points in a convex set. This yields the first exact integer optimization algorithm for general convex functions.

The work in this chapter is based in part on the paper [36] (joint with Chris Peikert and Santosh Vempala) as well as subsequent improvements.

7.1 Introduction

The Integer Linear Programming problem (ILP) is a classic NP-Complete problem that has received much attention within Computer Science and Operations Research. Given a rational polytope $P = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq \mathbf{b}\}$, the ILP is to decide whether P contains an integer point, i.e. whether $P \cap \mathbb{Z}^n \neq \emptyset$. More generally, for a convex set $K \subseteq \mathbb{R}^n$, the Convex Integer Programming feasibility problem (CIP) is to decide whether $K \cap \mathbb{Z}^n \neq \emptyset$. Both these problems are subclasses of the Integer Programming Problem (IP).

Starting with the classic cutting plane algorithm of Gomory [55], algorithms for ILP have been extensively studied over the last fifty years. Despite much effort, ILP remains one of the few NP-Complete problems for which no single exponential time algorithm is known. We note that when all the variables of the IP are binary (take values in $\{0, 1\}$), the trivial exhaustive search algorithm yields a straightforward 2^n time algorithm for an n -variable problem. Furthermore, using the standard assumption that 3-SAT is exponentially hard [107], the exhaustive search algorithm is essentially “optimal” for a general binary IP. When the variables are allowed to general integers however, developing an algorithm which yields any reasonable complexity bound is highly non-trivial. The first breakthrough algorithms in this area are due to Lenstra [84] and Kannan [70], where the former solves an n variable ILP in $2^{O(n^3)}$ time and the latter $2^{O(n)}n^{2.5n}$ time, and both use polynomial space. Subsequent to their work, many algorithms were discovered for solving more general problems such as counting the integer points in a rational polyhedron [14], parametric integer programming [71, 47], and integer programming over quasi-convex polynomials [63]. However, the core dimensional dependence of IP algorithms was not improved until recently [65]. In this work, Hildebrand and Köppe [65] develop a stronger deterministic ellipsoidal rounding scheme which allows them to reduce the dependence on n to $O(n)^{2n}$ while using $2^{O(n)}$ space. The central problem in this area thus remains: does there exist a $2^{O(n)}$ time algorithm for IP?

There are two main types of IP algorithms extant in the literature. The first are the Lenstra type algorithms, which use a “thinnest” direction of the feasible region to decompose the integer program into lower dimensional subproblems corresponding parallel lattice hyperplanes. The second are the Kannan type algorithms, which decompose the feasible region along lattice shifts of a linear subspace, generalizing the hyperplane decompositions of Lenstra.

The above algorithms all use lattice algorithms at their core, in particular algorithms for basis reduction and the shortest vector problem. Though the lattice problems solved within IP are inherently general norm problems, due to the lack of general norm techniques they were previously only solved approximately via reductions to ℓ_2 . In this chapter, we show how to obtain complexity improvements for IP using general norm techniques for lattice problems, as well as stronger structural results in the geometry of numbers.

7.1.1 Results

The fundamental tool for Lenstra type algorithms is Kinchine’s flatness theorem in the geometry of numbers (proved in the chapter 6). Recalling from the previous chapter, the flatness theorem states that

$$1 \leq \mu(K, \mathcal{L}) \cdot \lambda_1((K - K)^*, \mathcal{L}^*) \leq f(n),$$

The flatness theorem is most easily interpreted as follows: either K contains a lattice point in every translation ($\mu(K, L) \leq 1$), or there exist $\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}$ such that at most $\lfloor f(n) \rfloor + 1$ hyperplanes of the form $H_k = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{x} \rangle = k\}$, $k \in \mathbb{Z}$, intersect K (noting that any lattice point in K must be on one of these hyperplanes). Crucially, computing $\lambda_1((K - K)^*, \mathcal{L}^*)$ for a general convex body K exactly corresponds to a general norm SVP computation. The defining characteristic of Lenstra type algorithms is the use of “thin width” dual directions to decompose the IP along consecutive parallel hyperplanes (as described in the previous sentence).

Using the best known bounds on the flatness theorem, our algorithm achieves the following complexity:

Theorem 7.1.1 (IP-Lenstra). *Let K be a (\mathbf{a}_0, R) circumscribed convex set given by a strong separation oracle SEP_K . Let $\mathcal{L} \subseteq \mathbb{R}^n$ be an n -dimensional lattice given by a basis $B \in \mathbb{Q}^{n \times n}$. Then there is an algorithm which either decides that $K \cap \mathcal{L} = \emptyset$,*

or returns a point $\mathbf{y} \in K \cap \mathcal{L}$ in time $2^{O(n)}(n^{\frac{4}{3}} \text{polylog}(n))^n \text{poly}(\cdot)$ using $2^n \text{poly}(\cdot)$ space.

We note that any improvement on the bound $f(n)$ in the Flatness Theorem yields a corresponding improvement to the above algorithm. In general, we get an algorithm of complexity $2^{O(n)}f(n)^n$ (though an exact bound on $f(n)$ is required to run the algorithm). Hence, assuming the conjectured bound on $f(n)$, the complexity drops to $2^{O(n)}n^n$. However, given that $f(n) = \Omega(n)$ (even when $K = B_2^n$) it seems unlikely that Lenstra's approach can yield running times below $n^{\Omega(n)}$.

Our next result is an improved Kannan type algorithm for Integer Programming. The difference with respect to the Lenstra type algorithms is that the IP decompositions considered (i.e. how we decompose into subproblems) are general subspace decompositions as opposed to hyperplane decompositions.

In Lenstra's algorithm, the hyperplanes considered for K and lattice \mathcal{L} correspond to lattice hyperplanes orthogonal to a shortest non-zero vector \mathbf{y} of \mathcal{L}^* with respect to the norm induced by $(K - K)^*$. To generalize from this point of view, we formulate the decomposition in terms of fibers of a projection. Let $\pi_{\mathbf{y}}$ denote the orthogonal projection map onto $\text{span}(\mathbf{y})$. Now examine the set $A = \pi_{\mathbf{y}}(K) \cap \pi_{\mathbf{y}}(\mathcal{L})$. By the containment

$$\pi_{\mathbf{y}}(K \cap \mathcal{L}) \subseteq \pi_{\mathbf{y}}(K) \cap \pi_{\mathbf{y}}(\mathcal{L}) \tag{7.1.1}$$

we note that $K \cap \mathcal{L} \neq \emptyset \Rightarrow \pi_{\mathbf{y}}(K) \cap \pi_{\mathbf{y}}(\mathcal{L}) \neq \emptyset$. Since $\pi_{\mathbf{y}}$ is a 1-dimensional projection, we note that for each $\mathbf{a} \in A$, the fiber $\pi_{\mathbf{y}}^{-1}(\mathbf{a}) = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{v}, \mathbf{x} \rangle = \langle \mathbf{v}, \mathbf{a} \rangle\}$ corresponds to an $n - 1$ dimensional affine subspace (i.e. a hyperplane), satisfying $\pi_{\mathbf{y}}^{-1}(\mathbf{a}) \cap K \neq \emptyset$. From Equation (7.1.1) we have that

$$K \cap \mathcal{L} \subseteq \cup_{\mathbf{a} \in A} \pi_{\mathbf{y}}^{-1}(\mathbf{a}) \tag{7.1.2}$$

The above exactly recovers the hyperplane decomposition used in Lenstra's algorithm.

The idea behind Kannan’s algorithm is to use decompositions induced by the fibers of a general projection map, i.e. not just 1-dimensional as in Lenstra’s. In particular, for a general projection map π_W from \mathbb{R}^n to a k -dimensional linear subspace W , $1 \leq k \leq n$, letting $A = \pi_W(K) \cap \pi_W(\mathcal{L})$, we examine the subspace decomposition $K \cap \mathcal{L} \subseteq \cup_{\mathbf{a} \in A} \pi_W^{-1}(\mathbf{a})$.

There are two main difficulties with this approach. First, we need a general method for enumerating the subproblems, i.e. for computing the set $\pi_W(K) \cap \pi_W(\mathcal{L})$. When $\dim(W) = 1$ (as in Lenstra’s), this is straightforward, since $\pi_W(K)$ is an interval and $\pi_W(\mathcal{L})$ is generated by a single vector. For a general subspace W , computing $\pi_W(K) \cap \pi_W(\mathcal{L})$ corresponds to enumerating the lattice points (since $\pi_W(\mathcal{L})$ is a lattice as long as W is a lattice subspace of \mathcal{L}^*) inside a general convex set, which is non-trivial. To achieve this, we will rely on our lattice point enumeration algorithm (Algorithm 5.3). However we remember that algorithm Lattice-Enum only guarantees a running time proportional to $G(\pi_W(K), \pi_W(\mathcal{L}))$. Hence the second major problem is to find a projection π_W such that $G(\pi_W(K), \pi_W(\mathcal{L}))$ is “small”.

In Kannan’s paper [70], the above framework is directly implemented for the case where K is an ellipsoid. To obtain results for all convex bodies, Kannan first outer approximates the IP feasible region by an ellipsoid. All the lattice point enumeration and projection finding problems mentioned above are therefore reduced to their ℓ_2 counterparts. Our main contribution here is to give methods to solve these problems directly without incurring the losses due to ellipsoidal approximation. Combining these elements together allows us to give the following IP algorithm.

Theorem 7.1.2 (IP-Kannan). *Let $K \subseteq \mathbb{R}^n$ be a (\mathbf{a}_0, R) -circumscribed and let \mathcal{L} be an n -dimensional lattice with basis $B \in \mathbb{Q}^{n \times n}$. Then there exists an algorithm which either decides that $K \cap \mathcal{L} = \emptyset$ or outputs a point in $\mathbf{y} \in K \cap \mathcal{L}$ in time $O(n)^n \text{poly}(\cdot)$ using $2^n \text{poly}(\cdot)$ space.*

The complexity of the above algorithm is controlled, as in Lenstra’s algorithm,

by the number of subproblems created at each recursion node. As explained above, the subproblems correspond to the fibers of a projection map π_W , where the number of fibers is bounded by $G(\pi_W(K), \pi_W(\mathcal{L}))$. To build a “good” projection map π_W , we will use a construction of Kannan and Lovász [67] which yields a subspace W , $\dim(W) = k \in [n]$, satisfying $G(\pi_W(K), \pi_W(\mathcal{L})) = (3n)^k$ (see Lemma 7.4.4). Though we do not have control on the dimension of W , the stated bound combined with a simple recursion relation allows us to bound the maximum number of subproblems created during the algorithm by $2^{O(n)}n^n$ (which yields the dominant complexity term). Though this does not improve on the current conjectured runtime for Lenstra’s algorithm, it yields a pathway for improving on the current $n^{O(n)}$ runtime for IP. More precisely, a conjecture of Kannan and Lovász [67] essentially states that there always exists a subspace W , $\dim(W) = k$, such that $G(\pi_W(K), \pi(\mathcal{L})) = O(\log n)^k$. An algorithmic version of this conjecture (i.e. which computes the desired subspace W), would yield an $O(\log n)^n$ algorithm for IP (assuming it takes at most $O(\log n)^n$ time to find W), greatly improving upon the complexity of IP. We state this more formally in the following theorem:

Theorem 7.1.3 (IP-Kannan Extended, informal). *Assume that for any convex body $K \subseteq \mathbb{R}^n$ and n -dimensional lattice \mathcal{L} such that $\mu(K, \mathcal{L}) \geq 1$, there is a $2^{O(n)}g(n)^n$ time algorithm which computes a subspace $W \subseteq \mathbb{R}^n$, $\dim(W) = k \in [n]$, satisfying $G(\pi_W(K), \pi_W(\mathcal{L})) = 2^{O(k)}g(n)^r$. Then there is a $2^{O(n)}g(n)^n$ time algorithm for IP.*

A natural further question is whether the above IP feasibility solvers can be extended to work in the optimization setting. Our final result of this chapter shows that this can be achieved in a very general sense.

Theorem 7.1.4 (Convex IP). *Let $K \subseteq \mathbb{R}^n$ be a (\mathbf{a}_0, R) -circumscribed convex set, $f : K \rightarrow \mathbb{R}$ be a convex function equipped with subgradient oracle, and let \mathcal{L} denote an n -dimensional lattice with basis $B \in \mathbb{Q}^{n \times n}$. Then there exists a randomized algorithm*

which either decides that $K \cap \mathcal{L} = \emptyset$ or outputs a point $\mathbf{y} \in K \cap \mathcal{L}$ minimizing $f(\cdot)$ in expected time $O(n)^n \text{poly}(\cdot)$ using $2^n \text{poly}(\cdot)$ space.

We note that in the above theorem, if we are only interested in approximate minimizers (as opposed to exact), then there is a direct reduction to IP feasibility which uses a standard binary search on the objective (assuming it is easy to compute upper and lower bounds on the optimal objective value). Obtaining an exact minimizer using binary search however, seems to require us to make more assumptions about f than just convexity. Indeed, in most previous works on integer optimization, the exact description of f is used to bound the accuracy required for binary search to find an optimal solution. Our main contribution with the above result is to provide a method which completely avoids binary search, allowing us to handle general convex objective functions. As far as the author is aware, this is the first result of this type in the literature.

The following table summarizes all known complexity improvements for IP.

Constraints	Time	Space	Decomposition Type	Reference
Linear	$2^{O(n^3)}$	$\text{poly}(n)$	Hyperplane	[84]
Linear	$O(n)^{2.5n}$	$\text{poly}(n)$	Subspace	[70]
Quasi-Convex Polynomials	$O(n)^{2n}$	2^n	Hyperplane	[65]
General Convex	$\tilde{O}(n)^{\frac{4}{3}n}$	2^n	Hyperplane	[36]
General Convex	$O(n)^n$	2^n	Subspace	This thesis

Organization. The following sections are organized as follows. In the section 7.2, we show how to preprocess any integer program to one where the feasible region is well-sandwiched and the lattice basis is “short”. In sections 7.3 and 7.4 we give our implementations of Lenstra’s and Kannan’s algorithm respectively. In section 7.5, we give our algorithm for convex integer minimization.

7.2 Preprocessing the Integer Program

As a first step, we give a crucial preprocessing algorithm for an input IP which “normalizes” the initial feasible region and lattice basis. The preprocessing step allows us to deal with general convex sets which are not necessarily full dimensional. The reader will note that the time required to run the preprocessing is independent of “inner roundness” of K (say the size of the largest ball K contains in its affine hull). This is because if K is very “thin” the algorithm can use the discrete structure of the lattice to reduce dimension.

Lemma 7.2.1 (Algorithm IP-Preprocess). *Let $K \subseteq \mathbb{R}^n$ be a (\mathbf{a}_0, R) -circumscribed convex set given by a strong separation oracle SEP_K , and let \mathcal{L} denote an n -dimensional lattice given by a basis $B \in \mathbb{Q}^{n \times n}$, and let $H = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$, $A \in \mathbb{Q}^{m \times n}$, $\mathbf{b} \in \mathbb{Q}^m$, be an affine subspace. Then there is a $2^{O(n)}$ poly(\cdot) time algorithm which either decides that $K \cap \mathcal{L} \cap H = \emptyset$ or returns*

- (1) a shift $\mathbf{p} \in \mathcal{L}$,
- (2) a sublattice $\mathcal{L}' \subseteq \mathcal{L}$, $\dim(\mathcal{L}') = k \leq n$, given by a basis $\mathbf{b}'_1, \dots, \mathbf{b}'_k$,
- (3) a vector $\mathbf{a}'_0 \in \text{span}(\mathcal{L}')$ and radius $0 < R' \leq R$,
- (4) a convex set $K' = (K - \mathbf{p}) \cap (\mathbf{a}'_0 + R'B_2^n) \cap \text{span}(\mathcal{L}')$ (by its separation oracle),
- (5) an ellipsoid $E' = \{\mathbf{x} \in \text{span}(\mathcal{L}') : \mathbf{x}^t A' \mathbf{x} \leq 1\}$ and center $\mathbf{c}' \in \text{span}(\mathcal{L}')$

satisfying the following properties:

- (1) $K \cap \mathcal{L} \cap H = (K' \cap \mathcal{L}') + \mathbf{p}$.
- (2) $\mathbf{c}' + \frac{1}{\sqrt{k+1}(k+1)} E' \subseteq K' \subseteq \mathbf{c}' + E'$.
- (3) $\max_{1 \leq i \leq k} \|\mathbf{b}'_i\|_2 \leq 2\sqrt{k}R'$.
- (4) \mathbf{a}'_0 , \mathbf{c}' , A' , $\mathbf{b}'_1, \dots, \mathbf{b}'_k$, and \mathbf{p} have polynomial sized encodings.

Proof.

Basis Refinement: In this step, we use the circumscribing information for K and the affine space H to restrict \mathcal{L} to a shifted sublattice admitting a short basis.

First if $H = \emptyset$ return EMPTY. Else set $\bar{\mathbf{a}}_0 \leftarrow$ orthogonal projection of \mathbf{a}_0 onto H and $\bar{R} \leftarrow \sqrt{R^2 - \|\mathbf{a}_0 - \bar{\mathbf{a}}_0\|_2^2}$. Note that $\bar{R} \leq R$ and $K \cap H \subseteq \bar{\mathbf{a}}_0 + \bar{R}B_2^n$ and $\mathbf{a}_0 \in H$. Set $\mathbf{a}_0 \leftarrow \bar{\mathbf{a}}_0$ and $R \leftarrow \bar{R}$.

Use the MV algorithm (with some standard preprocessing) for CVP to compute a closest vector $\mathbf{p} \in \mathcal{L} \cap H$ to \mathbf{a}_0 in the ℓ_2 norm. Let $H_0 = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{0}\}$, and note that $H \cap \mathcal{L} = (\mathcal{L} \cap H_0) + \mathbf{p}$. If $\|\mathbf{p} - \mathbf{a}_0\|_2 > R$ return EMPTY (since $K \subseteq \mathbf{a}_0 + RB_2^n$). If not, set $K \leftarrow (K - \mathbf{p}) \cap H_0$ (which readily admits a separation oracle by Lemma 2.5.8) and $\mathbf{a}_0 \leftarrow \mathbf{a}_0 - \mathbf{p}$, noting now that $\mathbf{a}_0 \in H_0$ and $\|\mathbf{a}_0\| \leq R$.

Let $\mathcal{L} \leftarrow \mathcal{L} \cap H_0$ and compute a new basis $\mathbf{b}_1, \dots, \mathbf{b}_l$ where $l = \dim(H_0) \leq n$. Next we use MV algorithm again to compute linearly independent lattice vectors $\mathbf{v}_1, \dots, \mathbf{v}_l$ achieving the successive minima of \mathcal{L} , i.e. where $\|\mathbf{v}_i\|_2 = \lambda_i(\mathcal{L})$. Both invocations of the MV algorithm here take $2^{O(n)}$ poly(\cdot) time. Letting $\mathbf{v}_0 = \mathbf{0}$, perform the following procedure:

```

 $k \leftarrow l, \bar{\mathbf{a}}_0 \leftarrow \mathbf{a}_0, \bar{R} \leftarrow R.$ 
while  $\|\mathbf{v}_k\|_2 > 2\bar{R}$  do
   $k \leftarrow k - 1$ 
  if  $\|\mathbf{v}_k\|_2 \leq 2\bar{R}$  then
     $W_k \leftarrow \text{span}(\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_k).$ 
     $\bar{\mathbf{a}}_0 \leftarrow$  orthogonal projection of  $\mathbf{a}_0$  onto  $W_k.$ 
     $\bar{R} \leftarrow \sqrt{R^2 - \|\mathbf{a}_0 - \bar{\mathbf{a}}_0\|_2^2}.$ 

```

Let $\bar{\mathcal{L}} = \mathcal{L} \cap W_k$, where k is set to its final value in the above procedure. Let $\bar{\mathbf{a}}_0$ and \bar{R} be set to their final values as well.

Claim: $\mathcal{L} \cap (\mathbf{a}_0 + RB_2^n) \subseteq \bar{\mathcal{L}}.$

Proof. Take $\mathbf{y} \in \mathcal{L} \cap (\mathbf{a}_0 + RB_2^n)$. Let $i \geq 0$ be the minimum index such that $\mathbf{y} \in \mathcal{L} \cap W_i$,

where $W_i = \text{span}(\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_i)$ (note that i is well defined since $W_i = \text{span}(\mathcal{L})$). To prove the claim it suffices to show that $i \leq k$. Assume that $i > k$. For j , $0 \leq j \leq n$, let \mathbf{b}_j denote the orthogonal projection of \mathbf{a}_0 onto W_j and let $R_j = \sqrt{R^2 - \|\mathbf{a}_0 - \mathbf{b}_j\|_2^2}$. A straightforward computation reveals that

$$\mathbf{y} \in (\mathbf{a}_0 + RB_2^n) \cap W_j = (\mathbf{b}_j + R_j B_2^n) \cap W_j$$

As j increases, note that \mathbf{b}_j gets closer to \mathbf{a}_0 (since $W_j \subseteq W_{j+1}$), and hence $R_j \leq R_{j+1}$ for $0 \leq j \leq n-1$. By the above computations and by assumption on i , we see that $\mathbf{y} \in (\mathbf{a}_0 + RB_2^n) \cap W_i = (\mathbf{b}_i + R_i B_2^n) \cap W_i$, and hence $\|\mathbf{y} - \mathbf{b}_i\| \leq R_i$. Next since $\|\mathbf{a}_0\|_2 \leq R$, we have that $\mathbf{0} \in (\mathbf{a}_0 + RB_2^n) \cap W_i$, and therefore $\mathbf{0} \in \mathbf{b}_i + R_i B_2^n \Rightarrow \|\mathbf{b}_i\|_2 \leq R_i$. By our choice of i , we have that $\|\mathbf{y}\|_2 \geq \lambda_i(\mathcal{L}) = \|\mathbf{v}_i\|_2$. Since $i > k$, by the implementation of the procedure, we must have that $\|\mathbf{v}_i\|_2 > 2R_i$ and hence $\|\mathbf{y}\|_2 > 2R_i$. Therefore

$$\|\mathbf{y} - \mathbf{b}_i\|_2 \geq \|\mathbf{y}\|_2 - \|\mathbf{b}_i\|_2 > 2R_i - R_i = R_i$$

a clear contradiction. Therefore $i \leq k$ as needed. \square

Since $K \subseteq \mathbf{a}_0 + RB_2^n$, the above claim gives us that $K \cap \mathcal{L} = K \cap \bar{\mathcal{L}}$. Also, we remember that $K \cap W_k \subseteq (\bar{\mathbf{a}}_0 + \bar{R}B_2^n) \cap W_k$. Now set $\mathbf{a}_0 \leftarrow \bar{\mathbf{a}}_0$, $R \leftarrow \bar{R}$, and $\mathcal{L} \leftarrow \bar{\mathcal{L}}$ and $K \leftarrow K \cap W_k$.

Now using standard techniques, we may compute a basis $\mathbf{b}_1, \dots, \mathbf{b}_k$ for \mathcal{L} using $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_k$ satisfying $\max_i \|\mathbf{b}_i\|_2 \leq \sqrt{k} \|\mathbf{v}_k\|_2 \leq 2\sqrt{k}R$ in polynomial time.

After the above basis refinement, we have a (\mathbf{a}_0, R) -circumscribed convex set $K \subseteq W_k$, a lattice $\mathcal{L} \subseteq W_k$ with basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_k)$ satisfying $\max_i \|\mathbf{b}_i\|_2 \leq 2\sqrt{k}R$.

Localizing K : For the next step, we attempt to compute a strong enough ellipsoidal approximation of K . To do this, we use algorithm GLS-Round (Theorem 2.5.10), running against K (restricted to W_k) with parameter $\epsilon = \left(\frac{1}{4k}\right)^k \det(\mathcal{L})$ to

deterministically compute an ellipsoid $E + \mathbf{c} \subseteq W_k$ such that either (1) E sandwiches K well, i.e. $\mathbf{c} + \frac{1}{\sqrt{k+1}(k+1)}E \subseteq K \subseteq \mathbf{c} + E$, or (2) $\text{vol}_k(E) \leq \epsilon$ (i.e. E is tiny compared to the ‘sparsity’ of \mathcal{L}), or This step can be done in polynomial time. If we are in case (1), return E, \mathbf{c} as well as the current K, \mathcal{L} (and other associated parameters) as the preprocessing. If we are instead in case (2), proceed to the next part of the algorithm.

Reducing dimension: Here we are in case (2), and hence unable to compute a good ellipsoidal approximation of K . By the termination condition of the ellipsoidal rounding algorithm however, we know that the containing ellipsoid $E + \mathbf{c}$ for K has very small volume compared to $\det(\mathcal{L})$. We will use this information to reduce the effective dimension of K by one. To achieve this we will find a dual vector $\mathbf{y} \in \mathcal{L}^* \setminus \{0\}$, such that there exists at most one hyperplane of the form $H_s = \{\mathbf{x} \in W_k : \langle \mathbf{x}, \mathbf{y} \rangle = s\}$, $s \in \mathbb{Z}$, that intersects K . For $\mathbf{y} \in \mathcal{L}^* \setminus \{0\}$, we note any vector in $\mathbf{x} \in \mathcal{L}$ satisfies $\langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$, and hence every vector in $\mathcal{L} \cap K$ must lie on some hyperplane H_s , $s \in \mathbb{Z}$.

To find such a \mathbf{y} , we proceed as follows. If we are in case (1) above, we use the MV algorithm to compute a vector $\mathbf{y} \in \text{SVP}(E^*, \mathcal{L}^*)$, which can be done in $2^{O(k)} \text{poly}(\cdot)$ time. Noting that $(E - E)^* = \frac{1}{2}E^*$, we see that

$$\text{vol}_k((E - E)^*) = \left(\frac{1}{2}\right)^k \text{vol}_k(E^*) = \left(\frac{1}{2}\right)^k \text{vol}_k(B_2^k)^2 \frac{1}{\text{vol}_k(E)} > \left(\frac{1}{2k}\right)^k \frac{1}{\text{vol}_k(E)}$$

Given that $\text{vol}_k(E)^{\frac{1}{k}} \leq \epsilon = \frac{1}{4k} \det(\mathcal{L})^{\frac{1}{k}}$, from the above we see that

$$\text{vol}_k((E - E)^*)^{\frac{1}{k}} > \frac{1}{2k} \frac{1}{\text{vol}_k(E)^{\frac{1}{k}}} \geq 2 \frac{1}{\det(\mathcal{L})^{\frac{1}{k}}} = 2 \det(\mathcal{L}^*)^{\frac{1}{k}}$$

Since $(E - E)^* = \frac{1}{2}E^*$ is centrally symmetric, by Minkowski’s first theorem (Theorem 2.4.5) we have that $2\|\mathbf{y}\|_{E^*} = \|\mathbf{y}\|_{(E-E)^*} = \lambda_1((E - E)^*, \mathcal{L}^*) < 1$. We remember that $\|\mathbf{y}\|_{(E-E)^*} = \sup_{\mathbf{x} \in E} \langle \mathbf{y}, \mathbf{x} \rangle - \inf_{\mathbf{x} \in E} \langle \mathbf{y}, \mathbf{x} \rangle$ is the width of E with respect to \mathbf{y} . Since $\mathbf{y} \in \mathcal{L}^*$ we note that for any $\mathbf{x} \in (E + \mathbf{c}) \cap \mathcal{L}$, we must have that

$$\langle \mathbf{x}, \mathbf{y} \rangle \in (\langle \mathbf{y}, \mathbf{c} \rangle + [\inf_{\mathbf{x} \in E} \langle \mathbf{y}, \mathbf{x} \rangle, \sup_{\mathbf{x} \in E} \langle \mathbf{y}, \mathbf{x} \rangle]) \cap \mathbb{Z}.$$

Since E has width < 1 with respect to \mathbf{y} , it is easy to see that if $(E + \mathbf{c}) \cap \mathcal{L}$ is non-empty then all the lattice points in $(E + \mathbf{c}) \cap \mathcal{L}$ must lie on the hyperplane $H = \{\mathbf{x} \in W_k : \langle \mathbf{x}, \mathbf{y} \rangle = \lfloor \langle \mathbf{c}, \mathbf{y} \rangle \rfloor\}$. Since $K \subseteq E + \mathbf{c}$, it is also clearly the case that $K \cap \mathcal{L} \subseteq H \cap \mathcal{L}$.

If $H \cap (E + \mathbf{c}) = \emptyset$, return EMPTY. Else, call IP-Preprocess recursively on K, \mathcal{L} and H (with remaining parameters). If the recursive call returns EMPTY, return EMPTY. Else, letting $K', \mathcal{L}', \mathbf{p}'$ denote the returned convex set, lattice and shift, we return $K', \mathcal{L}', \mathbf{p}' + \mathbf{p}$ (along with remaining parameters). \square

7.3 An Improved Lenstra Type Algorithm

We are now ready to give our implementation of Lenstra's algorithm. Throughout the analysis, we shall let $f(n)$ the best current bound on the flatness constant for dimension n . We give present the outline of the algorithm (additional details will be provided in the proof).

Algorithm 7.1 IP-Lenstra(K, \mathcal{L}, H)

Input: (\mathbf{a}_0, R)-circumscribed convex set $K \subseteq \mathbb{R}^n$ with separation oracle SEP_K , n dimensional lattice \mathcal{L} with basis $B \in \mathbb{Q}^{n \times n}$, a rational affine subspace $H = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$.

Output: Return NULL if $K \cap \mathcal{L} \cap H = \emptyset$ or $\mathbf{y} \in K \cap \mathcal{L} \cap H$.

- 1: $(K, \mathcal{L}, \mathbf{p}) \leftarrow \text{IP-Preprocess}(K, \mathcal{L}, H)$.
 - 2: **if** $K = \emptyset$ **then return** NULL; **else if** $\mathbf{0} \in K$ **then return** \mathbf{p} .
 - 3: $\mathbf{y} \in \text{Shortest-Vectors}((K - K)^*, \mathcal{L}^*, \frac{1}{2})$.
 - 4: $\alpha = \min\{1, \frac{f(\dim(\mathcal{L}))}{\text{width}_K(\mathbf{y})-1}\}$. $K \leftarrow (1 - \alpha)\mathbf{a}_0 + \alpha K$.
 - 5: **for all** $s \in \{\langle \mathbf{y}, \mathbf{x} \rangle : \mathbf{x} \in K\} \cap \mathbb{Z}$ **do**
 - 6: $H_s \leftarrow \{\mathbf{x} \in \text{span}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle = s\}$.
 - 7: $\mathbf{y} \leftarrow \text{IP-Lenstra}(K, \mathcal{L}, H_s)$.
 - 8: **if** $\mathbf{y} \neq \text{NULL}$ **then return** $\mathbf{y} + \mathbf{p}$.
 - 9: **return** NULL.
-

Proof of Theorem 7.1.1.

IP ALGORITHM: To begin, we call IP-Preprocess(K, \mathcal{L}, H) algorithm which requires $2^{O(n)}$ poly(\cdot) time. If the preprocessing step concludes that $K \cap \mathcal{L} \cap H = \emptyset$,

terminate and return NULL. Otherwise, we recover a new lattice \mathcal{L} , $\dim(\mathcal{L}) = k$, with a short basis $\mathbf{b}_1, \dots, \mathbf{b}_k$, a well sandwiched convex set $K \subseteq \text{span}(\mathcal{L})$, $\mathbf{c} + \frac{1}{\sqrt{k+1(k+1)}}E \subseteq K \subseteq \mathbf{c} + E$, and a shift \mathbf{p} .

Branching on a “thinnest” direction: With the preprocessing phase complete, we will now attempt to reduce the problem of finding a vector in $K \cap \mathcal{L}$ to finding a lattice point in one of at most $f(k) + 1$ (noting that k is current dimension) hyperplane sections of K , i.e. reducing dimension by 1 at the cost of solving at most $f(k) + 1$ subproblems. The hyperplane sections in question will be of the form $H_s = \{\mathbf{x} \in \text{span}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle = s\}$, $s \in \mathbb{Z}$, for some $\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}$. Note that for a fixed $\mathbf{y} \in \mathcal{L}^*$, every $\mathbf{x} \in \mathcal{L}$ (and hence in $\mathcal{L} \cap K$) belongs to some H_s , namely $H_{\langle \mathbf{x}, \mathbf{y} \rangle}$ (since $\langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$).

To find a desired “thin” direction for K , we shall compute $\mathbf{y} \in \text{SVP}((K - K)^*, \mathcal{L}^*)$. To do this, we must build a weak distance oracle for $(K - K)^*$. Given that K is well sandwiched by E , using the Ellipsoid Method (theorem 2.5.9), for any $\mathbf{y} \in \mathcal{L}^*$ and $\epsilon > 0$, we may compute $l, u \in \mathbb{Q}$ satisfying

$$l - \frac{\epsilon}{2} \leq \inf_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle \leq l \quad u \leq \sup_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle \leq u + \frac{\epsilon}{2}$$

in polynomial time. We note now that

$$|\|\mathbf{y}\|_{(K-K)^*} - (u - l)| = \left| \sup_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle - \inf_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle - (u - l) \right| \leq \epsilon$$

as needed. Next, the SVP algorithm needs sandwiching guarantees on $(K - K)^*$. Given our guarantees on K , we see that $\frac{1}{2}E^* = (E - E)^* \subseteq K - K \subseteq \frac{1}{2}(k+1)\sqrt{k+1}E^*$. Technically, the algorithm Shortest-Vectors requires the sandwiching ratio with respect to euclidean balls, but this type of sandwiching is equivalent to ellipsoidal sandwiching after linear transformation. Having constructed a weak distance oracle for $(K - K)^*$ and computed the sandwiching guarantees, we may now call Shortest-Vectors($(K - K)^*, \mathcal{L}^*, \frac{1}{2}$) (Theorem 5.3.2) and retrieve $\mathbf{y} \in \mathcal{L}^* \setminus \{\mathbf{0}\}$ from the output.

Since the sandwiching guarantees are polynomial in k and the required accuracy is $\Omega(1)$, this can be executed in expected time $2^{O(n)}$ poly(\cdot) time.

Using the Ellipsoid Method (theorem 2.5.9) as above, we compute bounds $u, l \in \mathbb{Q}$ satisfying $u - \frac{1}{4} \leq \sup_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle \leq u$ and $l \leq \inf_{\mathbf{x} \in K} \langle \mathbf{y}, \mathbf{x} \rangle \leq l + \frac{1}{4}$ in polynomial time. By construction, we see that $u - l - \frac{1}{2} \leq \|\mathbf{y}\|_{(K-K)^*} \leq u - l$. Now by the guarantees on \mathbf{y} , we know that $\|\mathbf{y}\|_{(K-K)^*} \leq \lambda_1((K-K)^*, \mathcal{L}^*) + \frac{1}{2}$, and therefore

$$u - l - 1 \leq \|\mathbf{y}\|_{(K-K)^*} - \frac{1}{2} \leq \lambda_1((K-K)^*, \mathcal{L}^*) \leq \|\mathbf{y}\|_{(K-K)^*} \leq u - l \quad (7.3.1)$$

We will now distinguish two cases.

Case 1: $u - l - 1 < f(k)$. In this case, compute $A = [l, u] \cap \mathbb{Z}$, and let $H_s = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{x} \rangle = s\}$. By construction, for all $\mathbf{x} \in K$ we have that $\langle \mathbf{y}, \mathbf{x} \rangle \in [l, u]$, and furthermore if $\mathbf{x} \in \mathcal{L}$ then $\langle \mathbf{x}, \mathbf{y} \rangle = [l, u] \cap \mathbb{Z} = A$. Therefore $K \cap \mathcal{L} \subseteq \cup_{s \in A} H_s$. For each $s \in A$, we recursively call IP-Lenstra(K, \mathcal{L}, H_s). If for all $s \in A$, the recursive calls return NULL, return NULL. Else, if for some $s \in A$, the recursive call returns $\mathbf{y} \in K \cap \mathcal{L}$, return $\mathbf{y} + \mathbf{p}$. A straightforward computation reveals that the number of examined subproblems $|A|$ satisfies $|A| \leq \lfloor u - l \rfloor + 1 \leq \lfloor f(n) + 1 \rfloor + 1 = \lfloor f(k) \rfloor + 2$.

Case 2: $u - l - 1 \geq f(k)$. In this case, let $K' = (1 - \alpha)\mathbf{c} + \alpha K$ for $\alpha = \frac{f(k)}{u - l - 1}$. Note that a separation oracle for K' is readily available starting from a separation oracle for K (i.e. \mathbf{x} is separated from K' iff $\frac{\mathbf{x} - \alpha\mathbf{c}}{1 - \alpha}$ is separated from K). We shall rely on the following claim:

Claim: $K' \cap \mathcal{L} \neq \emptyset$, $K' \subseteq K$, and $K' \cap \mathcal{L} \subseteq \cup_{s \in A} H_s$ for $A = (\alpha[l, u] + (1 - \alpha)\langle \mathbf{y}, \mathbf{c} \rangle) \cap \mathbb{Z}$. Furthermore $|A| \leq \lfloor f(k) \rfloor + 2$.

Proof. First, since $0 \leq \alpha \leq 1$ and $\mathbf{c} \in K$, we clearly have that $K' = (1 - \alpha)\mathbf{c} + \alpha K \subseteq K$. Next, we note that

$$(K' - K')^* = ((1 - \alpha)\mathbf{c} + \alpha K - (1 - \alpha)\mathbf{c} - \alpha K)^* = (\alpha(K - K))^* = \frac{1}{\alpha}(K - K)^*.$$

By homogeneity

$$\lambda_1((K' - K')^*, \mathcal{L}^*) = \lambda_1\left(\frac{1}{\alpha}(K - K)^*, \mathcal{L}^*\right) = \alpha\lambda_1((K - K)^*, \mathcal{L}^*).$$

Hence by equation (7.3.1)

$$\lambda_1((K' - K')^*, \mathcal{L}^*) = \alpha\lambda_1((K - K)^*, \mathcal{L}^*) = \frac{f(k)}{u - l - 1}\lambda_1((K - K)^*, \mathcal{L}^*) \geq f(k)$$

By the Flatness Theorem, we must have that $\mu(K', \mathcal{L}) \leq 1$ and hence $K' \cap \mathcal{L} \neq \emptyset$.

Let $T = \{\langle \mathbf{y}, \mathbf{z} \rangle : \mathbf{z} \in K'\}$, and note that

$$\begin{aligned} T &= \{\langle \mathbf{y}, \mathbf{z} \rangle : \mathbf{z} \in \alpha K + (1 - \alpha)\mathbf{c}\} = \alpha\{\langle \mathbf{z}, \mathbf{y} \rangle : \mathbf{z} \in K\} + (1 - \alpha)\langle \mathbf{y}, \mathbf{c} \rangle \\ &\subseteq \alpha[l, u] + (1 - \alpha)\langle \mathbf{y}, \mathbf{c} \rangle \end{aligned}$$

For $x \in \mathcal{L}$, we know that $\langle x, y \rangle \in \mathbb{Z}$, hence if $x \in \mathcal{L} \cap K'$ we have that $\langle x, y \rangle \in T \cap \mathbb{Z} = A$. Therefore $K' \cap \mathcal{L} \subseteq \cup_{s \in A} H_s$. Now by same bound used in case 1, we see that $|A| \leq \lfloor \alpha(u - l) \rfloor + 1$. Noting that

$$\alpha(u - l) = \frac{f(k)}{u - l - 1}(u - l) = f(n) + \frac{f(k)}{u - l - 1} \leq f(k) + 1,$$

we get that $|A| \leq \lfloor f(k) + 1 \rfloor + 1 = \lfloor f(k) \rfloor + 2$, as needed. \square

By the claim, $K' \cap \mathcal{L} \neq \emptyset$, and hence it suffices to restrict our attention to K' and \mathcal{L} . Letting A be as in the claim, we recursively call $\text{IP-Lenstra}(K', \mathcal{L}, H_s)$ for each $s \in A$, and return the first lattice point found (which is guaranteed to exist since $K' \cap \mathcal{L} \neq \emptyset$).

RUNTIME: The correctness of the algorithm has already been discussed above, so it only remains to check that the runtime of the algorithm is bounded by $O(f(n))^n \text{poly}(\cdot)$. The algorithm above is recursive, where at each node of the recursion we perform the IP-Preprocess procedure and then break the problem into at most $\lfloor f(n) \rfloor + 2$ (since f is monotonic in n) subproblems which we solve recursively.

We now claim that the processing in each node of the recursion takes at most $2^{O(n)} \text{poly}(\cdot)$ time. Now we note that dimension decreases at each node, and the algorithms executed during node processing (IP-Preprocess, MV, GLS-Round, etc) have either polynomial or single exponential dependence on dimension. Therefore, the only thing we need to check is that the encoding length of the subspaces and bases used at each node (which the algorithms have polynomial dependence on) have encoding sizes bounded by a fixed polynomial in the original input. Here it is not hard to see, that the main concern lies with the sublattice bases used in the recursions nodes. We note that after each call to IP-Preprocess, the bases returned have length at most $2\sqrt{n}R$. Now we note that any lattice vector in $\mathcal{L} \subseteq \mathbb{R}^n$ of ℓ_2 norm $2\sqrt{n}R$ has encoding length bounded by $\text{poly}(\text{enc}\langle B \rangle, \log R, n)$ (where B is the top level basis). Hence immediately after each call to IP-Preprocess, the current basis indeed has encoding length bounded by a fixed polynomial in the input. Lastly, it is easy to see that in between calls to IP-Preprocess, the encoding size of all the computed parameters grows by at most a fixed polynomial as well. Hence, all computations during the execution of algorithm occur on inputs having polynomial encoding size as needed.

Lastly, since the branching factor of the recursion tree is $O(f(n))$ and the amount of processing at each recursion node is at most $2^{O(n)} \text{poly}(\cdot)$, we have that the total running time of the algorithm is at most $O(f(n))^n$ as needed. \square

7.4 An Improved Kannan Type Algorithm

To give our implementation of Kannan’s algorithm, we will first need to develop some tools in the geometry of numbers. In the next section, we give a volumetric characterization of the quality of a projection for lattice point enumeration. Following this, we will give an algorithmic version of a transference theorem of Kannan-Lovasz [67] which enables us to find a “good” (though not optimal) projection. Finally, in subsection

7.4.2 we give our implementation of the improved Kannan type algorithm.

7.4.1 Finding a Thin Projection

We first give a volumetric characterization (up to single exponential factor) of $G(K, \mathcal{L})$.

Lemma 7.4.1. *Let $K \subseteq \mathbb{R}^n$ be a convex body and let \mathcal{L} be an n -dimensional lattice with basis $B \in \mathbb{R}^{n \times n}$. Letting $C = K \pmod{B}$, we have that*

$$\frac{\text{vol}_n(K)^{\frac{1}{n}}}{\text{vol}_n(C)^{\frac{1}{n}}} \leq G(K, \mathcal{L})^{\frac{1}{n}} \leq \min\{2, 1 + \mu(K, \mathcal{L})\} \frac{\text{vol}_n(K)^{\frac{1}{n}}}{\text{vol}_n(C)^{\frac{1}{n}}} \quad (7.4.1)$$

Furthermore, if $\mu(K, \mathcal{L}) \leq 1$ then $\text{vol}(C) = \det(\mathcal{L})$.

Proof. Let $F = B[0, 1)^n$, i.e. the fundamental parallelepiped of \mathcal{L} with respect to B . We remember that the map $\mathbf{x} \rightarrow \mathbf{x} \pmod{B}$ sends \mathbb{R}^n to F by the rule $\mathbf{x} \pmod{B} = \mathbf{y} \in F$ satisfying $\mathbf{x} - \mathbf{y} \in \mathcal{L}$ (this uniquely determines \mathbf{y}). For the lower bound, we see that

$$\begin{aligned} \frac{\text{vol}_n(K)}{\text{vol}_n(C)} &= \frac{1}{\text{vol}_n(C)} \int_F |(\mathcal{L} + \mathbf{x}) \cap K| d\mathbf{x} = \frac{1}{\text{vol}_n(C)} \int_C |(\mathcal{L} + \mathbf{x}) \cap K| dx \\ &\leq \max_{\mathbf{x} \in C} |(\mathcal{L} + \mathbf{x}) \cap K| = G(K, \mathcal{L}) \end{aligned}$$

since for $\mathbf{x} \in F$, $|(\mathcal{L} + \mathbf{x}) \cap K| > 0$ iff $\mathbf{x} \in C$.

By shifting K , we may assume that $\mathbf{0} \in K$. Let $\alpha = \min\{1, \mu(K, \mathcal{L})\}$, and define

$$V = \{\mathbf{x} \in \alpha K : \mathbf{x} \leq_{lex} \mathbf{y}, \forall \mathbf{y} \in (\mathcal{L} + \mathbf{x}) \cap K\}$$

where \leq_{lex} is the standard lexicographic ordering on \mathbb{R}^n . Since \leq_{lex} is a total ordering, by construction of V we must have that $|\mathcal{L} + \mathbf{x} \cap V| \leq 1 \forall \mathbf{x} \in B[0, 1)^n$, and hence for distinct $\mathbf{x}, \mathbf{y} \in \mathcal{L}$ we get that $\mathbf{x} + V \cap \mathbf{y} + V = \emptyset$. Next since $\mathbf{0} \in K$ and $\alpha < 1$, $V \subseteq \alpha K \subseteq K$.

Claim: $\text{vol}_n(V) = \text{vol}_n(C)$

Proof. We first show that $V \pmod{B} = C$ and the map $\mathbf{x} \rightarrow \mathbf{x} \pmod{B}$ is injective on V . Since $\mathbf{x} \pmod{B} = \mathbf{y} \pmod{B}$ iff $\mathbf{x} - \mathbf{y} \in \mathcal{L}$, for distinct $\mathbf{x}, \mathbf{y} \in V$ we clearly have $\mathbf{x} - \mathbf{y} \notin \mathcal{L}$ and hence $\mathbf{x} \pmod{B} \neq \mathbf{y} \pmod{B}$. This proves injectivity on V .

Next we show that $V \pmod{B} = C'$ where $C' = \alpha K \pmod{B}$. If $\mathbf{w} \in C'$, then there exists $\mathbf{y} \in \alpha K$ such that $\mathbf{y} \pmod{B} = \mathbf{w}$. Let \mathbf{x} denote the lex least element of $(\mathcal{L} + \mathbf{y}) \cap \alpha K$. By construction $\mathbf{x} \in V$ and $\mathbf{x} - \mathbf{y} \in \mathcal{L}$, therefore $\mathbf{x} \pmod{B} = \mathbf{y} \pmod{B} = \mathbf{w}$ as needed.

If $\alpha = 1$, then clearly $C' = C$, hence $V \pmod{B} = C' = C$ as needed. If $\alpha < 1$, then $\alpha = \mu(K, \mathcal{L}) < 1$ and hence $|(\mathcal{L} + \mathbf{x}) \cap \alpha K| \geq 1$ for all $\mathbf{x} \in F$. Since $\alpha K \subseteq K$

$$F \subseteq \alpha K \pmod{B} \subseteq K \pmod{B} \subseteq F.$$

Therefore $V \pmod{B} = C' = C = F$ as needed. Since $V \pmod{B} = C$, by injectivity we get that

$$\begin{aligned} \text{vol}_n(C) &= \text{vol}_n(V \pmod{B}) = \text{vol}_n\left(\bigcup_{\mathbf{y} \in \mathcal{L}} V \cap (\mathbf{y} + F) \pmod{B}\right) \\ &= \text{vol}_n\left(\bigcup_{\mathbf{y} \in \mathcal{L}} V \cap (\mathbf{y} + F) - \mathbf{y}\right) = \sum_{\mathbf{y} \in \mathcal{L}} \text{vol}_n(V \cap (\mathbf{y} + F) - \mathbf{y}) \\ &= \sum_{\mathbf{y} \in \mathcal{L}} \text{vol}_n(V \cap (\mathbf{y} + F)) = \text{vol}_n(V) \end{aligned}$$

as needed. □

Take $\mathbf{x} \in \mathbb{R}^n$. We wish to bound $|K \cap (\mathcal{L} + \mathbf{x})|$. Since $V \subseteq \alpha K$, we note that for $\mathbf{y} \in K \cap (\mathcal{L} + \mathbf{x})$, we have that

$$\mathbf{y} + V \subseteq K + V \subseteq K + \alpha K = (1 + \alpha)K \tag{7.4.2}$$

Furthermore, since $\mathbf{x} + V \cap \mathbf{y} + V = \emptyset$ for distinct $\mathbf{x}, \mathbf{y} \in \mathcal{L}$ and $\text{vol}_n(V) = \text{vol}_n(C)$,

we have that

$$\begin{aligned} \text{vol}_n((1+\alpha)K) &\geq \text{vol}\left(\bigcup_{\mathbf{y} \in (\mathcal{L} + \mathbf{x}) \cap K} \mathbf{y} + V\right) = |(\mathcal{L} + \mathbf{x}) \cap K| \text{vol}_n(V) \\ &= |(\mathcal{L} + \mathbf{x}) \cap K| \text{vol}_n(C) \end{aligned}$$

Hence $|(\mathcal{L} + \mathbf{x}) \cap K| \leq \frac{\text{vol}_n((1+\alpha)K)}{\text{vol}_n(C)} = (1+\alpha)^n \frac{\text{vol}_n(K)}{\text{vol}_n(C)}$, as needed. Lastly, by the proof of the claim, if $\mu(K, \mathcal{L}) \leq 1$ then $\text{vol}_n(C) = \text{vol}_n(F) = \det(\mathcal{L})$ as needed. \square

Corollary 7.4.2. *Let $K \subseteq \mathbb{R}^n$ be a convex body with $\mu(K, \mathcal{L}) \leq 1$. Then for any linear subspace $W \subseteq \mathbb{R}^n$, $\dim(W) = k$, we have that*

$$\frac{\text{vol}_k(\pi_W(K))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq G(\pi_W(K), \pi_W(\mathcal{L}))^{\frac{1}{k}} \leq 2 \frac{\text{vol}_k(\pi_W(K))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}}$$

Proof. First, if $\pi_W(\mathcal{L})$ is not a lattice (i.e. is not discrete), then any open set A in W intersecting $\pi_W(\mathcal{L})$, i.e. for which $\pi_W(\mathcal{L}) \cap A \neq \emptyset$, satisfies $|\pi_W(\mathcal{L}) \cap A| = \infty$. Since $\pi_W(K)$ has non-empty interior, the previous statement implies that $G(\pi_W(K), \pi_W(\mathcal{L})) = \infty$. By convention, if $\pi_W(\mathcal{L})$ is not a lattice, we have that $\det(\pi_W(\mathcal{L})) = 0$, and hence both the lower and upper bounds above are infinite (again by convention) as needed.

If $\pi_W(\mathcal{L})$ is a lattice, I claim that $\mu(\pi_W(K), \pi_W(\mathcal{L})) \leq 1$. Since $\text{span}(\mathcal{L}) = W$, by definition we have that $\mu(\pi_W(K), \pi_W(\mathcal{L})) = \inf\{s \geq 0 : \pi_W(\mathcal{L}) + s\pi_W(K) = W\}$. Since $\mu(K, \mathcal{L}) \leq 1$, we have that $\mathcal{L} + K = \mathbb{R}^n$, and hence $\pi_W(\mathcal{L} + K) = \pi_W(\mathcal{L}) + \pi_W(K) = \pi_W(\mathbb{R}^n) = W$. Therefore $\mu(\pi_W(K), \pi_W(\mathcal{L})) \leq 1$.

Applying Lemma 7.4.1 to $\pi_W(K)$ and $\pi_W(\mathcal{L})$ (restricting to the subspace W) yields the result. \square

Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ denote a basis of \mathcal{L} , and let $\mathbf{b}_1^*, \dots, \mathbf{b}_n^*$ denote its gram-schmidt orthogonalization (see section 2.1.1 for a definition). Let π_i , $1 \leq i \leq n$, denote the orthogonal projection onto the orthogonal complement of $\text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{i-1})$.

A basis $\mathbf{b}_1, \dots, \mathbf{b}_n$ of \mathcal{L} is an HKZ basis with respect to $K - K$, if it satisfies the relations

$$\|\mathbf{b}_i^*\|_{\pi_i(K-K)} = \lambda_1(\pi_i(K - K), \pi_i(\mathcal{L})), \quad 1 \leq i \leq n$$

We now present the transference theorem of Kannan and Lovász [67] which relates crucial properties HKZ bases. This theorem will be at the center of our efficiency improvements for IP.

Theorem 7.4.3. *Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ denote a basis for \mathcal{L} satisfying, for some $\epsilon \in [0, 1]$,*

$$\|\mathbf{b}_i^*\|_{\pi_i(K-K)} \leq (1 + \epsilon)\lambda_1(\pi_i(K - K), \pi_i(\mathcal{L})), \quad i \in [n]. \quad ((1 + \epsilon) \text{ approximate HKZ basis})$$

Then

$$1 \leq \left(\sum_{i=1}^n \|\mathbf{b}_i^*\|_{\pi_i(K-K)} \right) \min_{1 \leq i \leq n} \frac{\text{vol}_{n-i+1}(\pi_i(K))^{\frac{1}{n-i+1}}}{\det(\pi_i(\mathcal{L}))^{\frac{1}{n-i+1}}} \leq (1 + \epsilon)n. \quad (7.4.3)$$

Furthermore, we have that

$$1 \leq \mu(K, \mathcal{L}) \inf_{\substack{W \subseteq \mathbb{R}^n \\ \dim(W)=k}} \frac{\text{vol}_k(\pi_W(K))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq n \quad (7.4.4)$$

where W ranges over all non-trivial linear subspaces of \mathbb{R}^n .

Proof. We prove the upper bound in equation (7.4.3). Let $j = \arg \max_{1 \leq i \leq n} \|\mathbf{b}_i^*\|_{\pi_i(K-K)}$.

Clearly

$$\sum_{i=1}^n \|\mathbf{b}_i^*\|_{\pi_i(K-K)} \leq n \|\mathbf{b}_j^*\|_{\pi_j(K-K)} \quad (7.4.5)$$

By Minkowski's first theorem, we have that

$$\begin{aligned} \|\mathbf{b}_j^*\|_{\pi_j(K-K)} &= (1 + \epsilon)\lambda_1(\pi_j(K - K), \pi_j(\mathcal{L})) \\ &\leq 2(1 + \epsilon) \frac{\det(\pi_j(\mathcal{L}))^{\frac{1}{n-j+1}}}{\text{vol}_{n-j+1}(\pi_j(K - K))^{\frac{1}{n-j+1}}} \leq (1 + \epsilon) \frac{\det(\pi_j(\mathcal{L}))^{\frac{1}{n-j+1}}}{\text{vol}_{n-j+1}(\pi_j(K))^{\frac{1}{n-j+1}}} \end{aligned} \quad (7.4.6)$$

where the last inequality follows by the Brunn-Minkowski inequality, i.e.

$$\begin{aligned} \text{vol}_{n-j+1}(\pi_j(K - K))^{\frac{1}{n-j+1}} &\geq \text{vol}_{n-j+1}(\pi_j(K))^{\frac{1}{n-j+1}} + \text{vol}_{n-j+1}(\pi_j(-K))^{\frac{1}{n-j+1}} \\ &= 2\text{vol}_{n-j+1}(\pi_j(K))^{\frac{1}{n-j+1}} \end{aligned}$$

The bound now follows by combining equations (7.4.5) and (7.4.6).

We prove the upper bound for equation (7.4.4). By Lemma 2.4.7, we have that

$$\mu(K, \mathcal{L}) \leq \sum_{i=1}^n \|\mathbf{b}_i^*\|_{\pi_i(K-K)}$$

Furthermore, by direct inclusion we have that

$$\inf_{\substack{W \subseteq \mathbb{R}^n \\ \dim(W)=k}} \frac{\text{vol}_k(\pi_W(K))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq \min_{1 \leq i \leq n} \frac{\text{vol}_{n-i+1}(\pi_i(K))^{\frac{1}{n-i+1}}}{\det(\pi_i(\mathcal{L}))^{\frac{1}{n-i+1}}}$$

Therefore the upper bound on equation (7.4.4) is direct consequence of the upper bound on equation (7.4.3) (setting $\epsilon = 0$). From this we also have that the lower bound on equation (7.4.4) implies the lower bound for equation (7.4.3)

We prove the lower bound for equation (7.4.4). For a subspace $W \subseteq \mathbb{R}^n$, $\dim(W) = k$, we note by homogeneity that $\mu(K, \mathcal{L}) \text{vol}_k(\pi_W(K))^{\frac{1}{k}} = \text{vol}_k(\pi_W(\mu(K, \mathcal{L})K))^{\frac{1}{k}}$. Hence we may assume that $\mu(K, \mathcal{L}) = 1$.

It now suffices to show that $\text{vol}_k(\pi_W(K)) \geq \det(\pi_W(\mathcal{L}))$. From the proof of Corollary 7.4.2, we have that $\mu(\pi_W(K), \pi_W(\mathcal{L})) \leq \mu(K, \mathcal{L}) = 1$. Therefore, $\pi_W(K)$ contains a fundamental domain of $\pi_W(\mathcal{L})$ and hence $\text{vol}_k(\pi_W(K)) \geq \det(\pi_W(\mathcal{L}))$ as needed. \square

The next algorithm essentially makes the above transference theorem algorithmic and reinterprets its conclusion (using corollary 7.4.2) in terms of $G(K, \mathcal{L})$ instead of volumes.

Lemma 7.4.4 (Algorithm Find-Subspace). *Let $K \subseteq \mathbb{R}^n$ be an (\mathbf{a}_0, r, R) -centered convex set with separation oracle SEP_K , and let \mathcal{L} be an n -dimensional lattice with basis $B \in \mathbb{Q}^{n \times n}$. There there exists a $2^{O(n)}$ poly(\cdot) time and 2^n poly(\cdot) space algorithm which returns a convex set $K' \subseteq K$ (by its separation oracle) and a subspace $W \subseteq \mathbb{R}^n$, $\dim(W) = k$, $1 \leq k \leq n$, such that*

$$K' \cap \mathcal{L} \neq \emptyset \Leftrightarrow K \cap \mathcal{L} \neq \emptyset$$

and

$$G(\pi_W(K'), \pi_W(\mathcal{L}))^{\frac{1}{k}} \leq 3n$$

Proof. Since K is centered and equipped with a separation oracle SEP_K , for any projection $\pi : \mathbb{R}^n \rightarrow W$, we can implement a weak distance oracle $D_{\pi(K-K)}$ for $\|\cdot\|_{\pi(K-K)}$, where $D_{\pi(K-K)}(\mathbf{x}, \epsilon)$ executes in polynomial time.

We begin by building a $(1 + \frac{1}{24n})$ -approximate HKZ basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ for \mathcal{L} as follows:

$\pi_1 \leftarrow$ identity on \mathbb{R}^n

for $i \in 1 \dots n$ **do**

 Build distance oracle for $\pi_i(K - K)$ and basis for $\pi_i(\mathcal{L})$.

 Compute $\mathbf{v}_i^* \in \text{Shortest-Vectors}(\pi_i(K - K), \pi_i(\mathcal{L}), \frac{1}{24n})$.

 Compute $\mathbf{v}_i \in \mathcal{L}$ such that $\pi_i(\mathbf{v}_i) = \mathbf{v}_i^*$.

$\pi_{i+1} \leftarrow$ orthogonal projection map onto $\text{span}(\mathbf{v}_1, \dots, \mathbf{v}_i)^\perp$.

From the guarantees on $\text{Shortest-Vectors}(\pi_i(K - K), \pi_i(\mathcal{L}), \frac{1}{24n})$, it is clear that the algorithm outputs a basis satisfying $\|\mathbf{v}_i^*\|_{\pi_i(K-K)} \leq (1 + \frac{1}{24n})\lambda_1(\pi_i(K - K), \pi_i(\mathcal{L}))$.

We now bound its runtime. Each call to the Shortest-Vectors procedure executes in $2^{O(n)} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space, and the runtime of the remaining loop operations is polynomial. Since the loop executes n times, we see that the entire runtime is bounded by $2^{O(n)} \text{poly}(\cdot)$ as needed.

Let $t = \sum_{i=1}^n \frac{24n}{24n-1} D_{\pi_i(K-K)}(\mathbf{v}_i^*, \frac{1}{24n})$ and let $\alpha = \min\{1, t\}$. First note that from the guarantees on the distance oracles, we have that

$$(1 - \frac{1}{24n}) \sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)} \leq \sum_{i=1}^n D_{\pi_i(K-K)}(\mathbf{v}_i^*, \frac{1}{24n}) \leq (1 + \frac{1}{24n}) \sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)}$$

Multiplying through by $\frac{24n}{24n-1}$, we get that

$$\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)} \leq t \leq \frac{24n+1}{24n-1} \sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)} \quad (7.4.7)$$

Let $K' = (1 - \alpha)\mathbf{a}_0 + \alpha K$. We claim that $K' \cap \mathcal{L} \neq \emptyset \Leftrightarrow K \cap \mathcal{L} = \emptyset$. If $\alpha = 1$, then the statement is trivially true since $K' = K$. If $\alpha < 1$ then $\alpha = t < 1$. Since

$K' - K' = t(K - K)$, we see that

$$\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(t(K-K))} = \frac{1}{t} \sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)}$$

Therefore, using equation (7.4.7) we get that

$$\frac{24n-1}{24n+1} \leq \sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(t(K-K))} \leq 1 \quad (7.4.8)$$

Now by Lemma 2.4.7, we have that

$$\mu(K', \mathcal{L}) = \mu(tK, \mathcal{L}) \leq \sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(t(K-K))} \leq 1$$

Therefore $K' \cap \mathcal{L} \neq \emptyset$. Since $0 \leq \alpha \leq 1$ and $\mathbf{a}_0 \in K$, we have that $K' \subseteq K$ and hence $K \cap \mathcal{L} \supseteq K' \cap \mathcal{L} \neq \emptyset$ as needed.

Now compute $j = \arg \max_{1 \leq i \leq n} D_{\pi_i(K-K)}(\mathbf{v}_i^*, \frac{1}{24n})$. Let $W = \text{span}(\mathbf{v}_1, \dots, \mathbf{v}_{j-1})^\perp$ and $k = n - j + 1$, where we note that $\pi_j = \pi_W$. By the guarantees on the distance oracles, for $i \in [n]$, we see that

$$(1 + \frac{1}{24n}) \|\mathbf{v}_j^*\|_{\pi_j(K-K)} \geq D_{\pi_j(K-K)}(\mathbf{v}_j^*, \frac{1}{24n}) \geq D_{\pi_i(K-K)}(\mathbf{v}_i^*, \frac{1}{24n}) \geq (1 - \frac{1}{24n}) \|\mathbf{v}_i^*\|_{\pi_i(K-K)}$$

From the above, we get that $(\frac{24n+1}{24n-1}) \|\mathbf{v}_j^*\|_{\pi_j(K-K)} \geq \max_{i \in [k]} \|\mathbf{v}_i^*\|_{\pi_i(K-K)}$, and hence

$$\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)} \leq n \left(\frac{24n+1}{24n-1} \right) \|\mathbf{v}_j^*\|_{\pi_j(K-K)} \quad (7.4.9)$$

The algorithm now returns K' and W . The following claim establishes the correctness of the algorithm.

Claim: $G(\pi_W(K'), \pi_W(\mathcal{L}))^{\frac{1}{k}} \leq 3n$.

By construction, we note that $t \leq \alpha$ and hence $\frac{t}{\alpha} \geq 1$. Therefore

$$G(\pi_W(K'), \pi_W(\mathcal{L})) \leq G(\pi_W(\frac{t}{\alpha}K'), \pi_W(\mathcal{L})) = G(\pi_W(tK), \pi_W(\mathcal{L}))$$

Hence it suffices to bound $G(\pi_W(tK), \pi_W(\mathcal{L}))$. By equation (7.4.8), we know that $\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(t(K-K))} \leq 1$. Therefore by Lemma 2.4.7, we have that $\mu(tK, \mathcal{L}) \leq 1$. Now

since $\|\mathbf{v}_j^*\|_{\pi_j(K-K)} \leq \left(\frac{24n+1}{24n}\right)\lambda_1(\pi_j(K-K), \pi_j(\mathcal{L}))$ and $\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(K-K)} \leq n\left(\frac{24n+1}{24n-1}\right)$, by our choice of W and the proof of Theorem 7.4.3, we have that

$$\left(\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(t(K-K))}\right) \frac{\text{vol}_k(\pi_W(tK))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq n \left(\frac{24n+1}{24n}\right) \left(\frac{24n+1}{24n-1}\right)$$

Combining the above with lower bound $\sum_{i=1}^n \|\mathbf{v}_i^*\|_{\pi_i(t(K-K))} \geq \frac{24n-1}{24n+1}$ from Equation 7.4.8, we get that

$$\frac{\text{vol}_k(\pi_W(tK))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq n \left(\frac{24n+1}{24n}\right) \left(\frac{24n+1}{24n-1}\right)^2 \leq \frac{5}{4}n$$

for $n \geq 1$. Since $\mu(tK, \mathcal{L}) \leq 1$ by Corollary 7.4.2 we have that

$$G(\pi_W(tK), \pi_W(\mathcal{L})) \leq 2 \frac{\text{vol}_k(\pi_W(tK))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq 2\left(\frac{5}{4}n\right) \leq 3n$$

as needed. □

As we will see from the our new IP algorithm's analysis, an improvement on the $G(\pi_W(K), \pi_W(\mathcal{L}))$ bound for the returned subspace in the above algorithm will immediately translate to an improvement in the complexity of IP. Though the limit for the complexity of Lenstra type algorithms seems likely to be $n^{\Omega(n)}$, the subspace decomposition approach we describe may enable us to do much better. In particular, it is possible that using just subspace decomposition techniques one maybe able to develop an $O(\log n)^n$ time algorithm for IP. This hope is encoded in the following conjecture of Kannan and Lovász [67].

Conjecture 7.4.5 (Subspace Flatness Conjecture). For a convex body $K \subseteq \mathbb{R}^n$ and n dimensional lattice \mathcal{L} , we have that

$$1 \leq \mu(K, \mathcal{L}) \inf_{\substack{W \subseteq \mathbb{R}^n \\ \dim(W)=k}} \frac{\text{vol}_k(\pi_W(K))^{\frac{1}{k}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{k}}} \leq O(\log n) \tag{7.4.10}$$

where W ranges over all non-trivial linear subspaces of \mathbb{R}^n .

In the next section, we will see that an algorithmic version of the above conjecture (or any progress towards it) will yield a faster algorithm for IP.

7.4.2 The Improved Algorithm

We are now ready to give the proof for the new IP algorithm (see Algorithm 7.2 for the description). Note that if no subspace H is provided for input $K, \mathcal{L} \subseteq \mathbb{R}^n$, we simply call $\text{IP-Kannan}(K, \mathcal{L}, \mathbb{R}^n)$.

Algorithm 7.2 $\text{IP-Kannan}(K, \mathcal{L}, H)$

Input: (\mathbf{a}_0, R) -circumscribed convex set $K \subseteq \mathbb{R}^n$ with separation oracle SEP_K , n dimensional lattice \mathcal{L} with basis $B \in \mathbb{Q}^{n \times n}$, a rational affine subspace $H = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$.

Output: Return NULL if $K \cap \mathcal{L} \cap H = \emptyset$ or $\mathbf{y} \in K \cap \mathcal{L} \cap H$.

- 1: $(K, \mathcal{L}, \mathbf{p}) \leftarrow \text{IP-Preprocess}(K, \mathcal{L}, H)$.
 - 2: **if** $K = \emptyset$ **then return** NULL; **else if** $\mathbf{0} \in K$ **then return** \mathbf{p} .
 - 3: $(K, W) \leftarrow \text{Find-Subspace}(K, \mathcal{L})$.
 - 4: **for all** $\mathbf{s} \in \text{Lattice-Enum}(\pi_W(K), \pi_W(\mathcal{L}), r)$ **do**
 - 5: $\mathbf{y} \leftarrow \text{IP-Kannan}(K, \mathcal{L}, \pi_W^{-1}(\mathbf{s}))$.
 - 6: **if** $\mathbf{y} \neq \text{NULL}$ **then return** $\mathbf{y} + \mathbf{p}$.
 - 7: **return** NULL.
-

Proof of Theorem 7.1.2.

Correctness: As most of the processing occurs in previously analyzed subroutines, the correctness of the IP-Kannan algorithm is straightforward. We first call IP-Preprocess on K, \mathcal{L} and H to get an equivalent and “nicer” integer program. Next, we call Find-Subspace on the preprocessed K and \mathcal{L} . Find-Subspace possibly rescales K (if K is too “wide”), which is guaranteed contain a lattice point if the original K contained one, and then finds a subspace W such that $G(\pi_W(K), \pi_W(\mathcal{L}))$ is “small”. Next, we use the Lattice-Enum algorithm to decompose the IP along lattice shifts of the subspace W^\perp , and recurse on each subproblem.

Runtime: Given the correctness, we now argue the runtime. Using essentially the same analysis as in IP-Lenstra, we have that have the encoding sizes of all bases and subspaces in the recursion nodes are polynomial in the top level inputs. Therefore, at each node, we have that calls to IP-Preprocess and Find-Subspace all take at most

$2^{O(n)}$ poly(\cdot) time. The time to compute the subspace decompositions however varies at each node, so we require an aggregate analysis.

Examine a recursive call to IP-Kannan on a convex set K, \mathcal{L} and subspace H , where $\min\{\dim(H), \dim(K), \dim(\mathcal{L})\} \leq n - d$ (at the top level node $d = 0$). Notice that after IP-Preprocess, we are left with a new convex K and lattice \mathcal{L} satisfying $\dim(K), \dim(\mathcal{L}) \leq n - d$. Therefore, after running Find-Subspace(K, \mathcal{L}), we are left with a rescaled convex set K , and subspace W , $\dim(W) = l \leq n - d$, such that $G(\pi_W(K), \pi_W(\mathcal{L})) \leq (3(n - d))^l$. Hence, the call to Lattice-Enum($\pi_W(K), \pi_W(\mathcal{L}), r$) outputs at most $(3C_1(n - d))^l$ points in time $(3C_2(n - d))^l$ poly(\cdot) (where $C_2 \geq C_1$) using 2^l poly(\cdot) space. From the perspective of the worst case analysis, we may assume that $\dim(K) = \dim(\mathcal{L}) = n - d$ and that exactly $(3C_1(n - d))^l$ points are enumerated, each corresponding to a subproblem of dimension at $n - d - l$. In this case, we may charge $(C_2/C_1)^l$ poly(\cdot) = $2^{O(n)}$ poly(\cdot) to each created recursion node to account for the time needed enumerate $\pi_W(K) \cap \pi_W(\mathcal{L})$.

From the above analysis, if we let $T(n)$ denote the maximum size of the recursion tree for an integer program of dimension n , then the total running time is bounded by $2^{O(n)}T(n)$ poly(\cdot). Clearly $T(0) = 1$, and from the above analysis, $T(n)$ satisfies the recursion relation

$$T(n) \leq 1 + \max_{1 \leq i \leq n} (3C_1 n)^i T(n - i)$$

I claim that $T(n) \leq (3C_1(n + 1))^n$. The base case holds since $T(0) = 1 = (3C_1)^0$.

Next, we have that

$$\begin{aligned} T(n) &\leq 1 + \max_{i \leq i \leq n} (3C_1 n)^i T(n - i) \leq 1 + \max_{i \leq i \leq n} (3C_1 n)^i (3C_1(n - i + 1))^{n-i} \\ &= 1 + (3C_1)^n \max_{i \leq i \leq n} n^i (n - i + 1)^{n-i} \end{aligned}$$

Now by the AM-GM inequality, we have that

$$\max_{1 \leq i \leq n} n^i (n - i + 1)^{n-i} \leq \max_{1 \leq i \leq n} \left(\frac{in + (n - i)(n - i + 1)}{n} \right)^n = n^n$$

where the last equality holds since $\max_{1 \leq i \leq n} \frac{i+(n-i)(n-i+1)}{n} = n$, where the max is attained for $i = 1, n$. Putting it all together, we get that

$$T(n) \leq 1 + (3C_1)^n \max_{i \leq i \leq n} n^i (n - i + 1)^{n-i} \leq 1 + (3C_1 n)^n \leq (3C_1(n + 1))^n$$

as needed. Therefore the total running time is bounded by $2^{O(n)}T(n) \text{poly}(\cdot) = 2^{O(n)}n^n \text{poly}(\cdot)$ as needed. \square

We now discuss a possible pathway for improvement based on the above algorithmic approach. A major conjectured source of improvement lies in the possibility of a better subspace finding algorithm. The current Find-Subspace algorithm is capable of finding a subspace inducing at most $(3n)^r$ subproblems of dimension $n - r$, for some $r \in [n]$. However, as suggested by Conjecture 7.4.5, it seems possible that an algorithm exists which returns a subspace inducing at most $O(\log n)^r$ subproblems of dimension $n - r$, drastically decreasing the total size of the recursion tree.

The following theorem relates the potential complexity improvements for IP given an improved projection finding algorithm. To be able to make minimal assumptions on the projection finding procedure, we will only run it when the covering radius of the input body is already $\Omega(1)$. This means we will avoid the scaling technique used by the above algorithm (as well as in our implementation of Lenstra's algorithm), and instead rely on the near central lattice point finder from Section 5.6. The only drawback to this alternate approach is that the implied algorithm will be randomized (though it remains Las Vegas). The exact theorem is stated as follows.

Theorem 7.4.6 (IP-Kannan Extension). *Let $g : \mathbb{N} \rightarrow \mathbb{N}$ denote a non-decreasing function. Assume there exists an algorithm Thin-Projection that given a centered convex body $K \subseteq \mathbb{R}^n$ and n dimensional lattice \mathcal{L} satisfying $\mu(K, \mathcal{L}) \geq \frac{1}{2}$, outputs a linear subspace $W \subseteq \mathbb{R}^n$, $\dim(W) = k \in [n]$, such that $G(\pi_W(K), \pi_W(\mathcal{L})) \leq g(n)^n$ in time $2^{O(n)}g(n)^n \text{poly}(\cdot)$. Then there is a randomized algorithm which solves the Convex Integer Programming feasibility problem in expected time $2^{O(n)}g(n)^n \text{poly}(\cdot)$*

time. Furthermore, the space complexity of the algorithm (notwithstanding the space needed to run *Thin-Projection*) is $2^n \text{poly}(\cdot)$.

We note that the current projection finding algorithm produces a subspace satisfying the conditions of the above theorem with $g(n) = 3n$, and hence yields an $2^{O(n)}n^n$ time algorithm for IP. Assuming Conjecture 7.4.5, it maybe possible to develop a projection finding algorithm satisfying $g(n) = O(\log n)$ above, yielding a $O(\log n)^n$ time algorithm for IP.

The guarantees of the above theorem are realized by the following algorithm.

Algorithm 7.3 IP-Kannan-Ext(K, \mathcal{L}, H)

Input: (\mathbf{a}_0, R) -circumscribed convex set $K \subseteq \mathbb{R}^n$ with separation oracle SEP_K , n dimensional lattice \mathcal{L} with basis $B \in \mathbb{Q}^{n \times n}$, a rational affine subspace $H = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$.

Output: Return NULL if $K \cap \mathcal{L} \cap H = \emptyset$ or $\mathbf{y} \in K \cap \mathcal{L} \cap H$.

- 1: $(K, \mathcal{L}, \mathbf{p}) \leftarrow \text{IP-Preprocess}(K, \mathcal{L}, H)$.
 - 2: **if** $K = \emptyset$ **then return** NULL; **else if** $\mathbf{0} \in K$ **then return** \mathbf{p} .
 - 3: $(\mathbf{c}, \mathbf{y}) \leftarrow \text{Central-Lat-Pt}(K, \mathcal{L}, 1)$.
 - 4: **if** $\mathbf{y} \in K$ **then return** $\mathbf{y} + \mathbf{p}$.
 - 5: $W \leftarrow \text{Thin-Projection}(K, \mathcal{L})$.
 - 6: **for all** $\mathbf{s} \in \text{Lattice-Enum}(\pi_W(K), \pi_W(\mathcal{L}), r)$ **do**
 - 7: $\mathbf{y} \leftarrow \text{IP-Kannan-Ext}(K, \mathcal{L}, \pi_W^{-1}(\mathbf{s}))$.
 - 8: **if** $\mathbf{y} \neq \text{NULL}$ **then return** $\mathbf{y} + \mathbf{p}$.
 - 9: **return** NULL.
-

Analysis of IP-Kannan-Ext.

Correctness: The correctness of the algorithm follows almost the same line of reasoning as Algorithm 7.2 (Kannan-IP). The main difference is the use of the Central-Lat-Pt algorithm, which replaces the scaling technique of IP-Kannan. We note that we only call Central-Lat-Pt on K, \mathcal{L} after preprocessing, and hence K is well-centered as needed by the algorithm. We now show that after running Central-Lat-Pt, either we find a feasible lattice point or we guarantee that $\mu(K, \mathcal{L}) \geq \frac{1}{2}$.

Let (\mathbf{c}, \mathbf{y}) denote the center and lattice point returned by Central-Lat-Pt($K, \mathcal{L}, 1$). If $\mathbf{y} \in K$, then by the preprocessing guarantee, $\mathbf{y} + \mathbf{p}$ is a feasible lattice point for

the original problem. Otherwise, if $\mathbf{y} \notin K$, by the guarantees of Central-Lat-Pt on $K, \mathcal{L}, 1$ we have that

$$1 \leq \|\mathbf{y} - \mathbf{c}\|_{K-\mathbf{c}} \leq 2d_{K-\mathbf{c}}(\mathcal{L}, \mathbf{c})$$

The above implies that $d_{K-\mathbf{c}}(\mathcal{L}, \mathbf{c}) \geq \frac{1}{2}$ and hence the body $K' = \frac{1}{2}K + \frac{1}{2}\mathbf{c}$ can only contain lattice points on its boundary. In turn, this shows that $1 \geq \mu(K', \mathcal{L}) = 2\mu(K, \mathcal{L})$, and hence $\mu(K, \mathcal{L}) \geq \frac{1}{2}$ as needed.

If the call to Central-Lat-Pt does not produce a feasible lattice point, the algorithm proceeds to decompose the feasible region into subproblems using the projection π_W (where W is produced by algorithm Thin-Projection). The correctness of this step follows exactly in the same way as for IP-Kannan. The algorithm is thus correct.

Runtime: The algorithm begins with the standard preprocessing (IP-Preprocess) and then proceeds with a call to Central-Lat-Pt. Together these calls require at most $2^{O(n)}$ poly(\cdot) expected time and 2^n poly(\cdot) space. If the algorithm does not find a feasible lattice point after the call to Central-Lat-Pt, it continues by executing the subroutine Thin-Projection on K and \mathcal{L} . Let $l = \dim(K) = \dim(\mathcal{L}) \leq n$. By the proof of correctness, we are guaranteed that $\mu(K, \mathcal{L}) \geq \frac{1}{2}$ before the call to Thin-Projection. Therefore by the guarantees on Thin-Projection, in $2^{O(l)}g(l)^l$ poly(\cdot) time it returns a subspace W , $\dim(W) = k \in [l]$, satisfying $G(\pi_W(K), \pi_W(\mathcal{L})) \leq g(l)^k$. From here, the algorithm enumerates the $l - k$ -dimensional subproblems indexed by $\pi_W(K) \cap \pi_W(\mathcal{L})$, where we note that $|\pi_W(K) \cap \pi_W(\mathcal{L})| \leq g(l)^k \leq g(n)^k$ (since g is non-decreasing). By the guarantees on algorithm Lattice-Enum, this requires $2^{O(k)}g(l)^k$ poly(\cdot) time and 2^l poly(\cdot) space. Lastly, the algorithm recurses on each of these $l - k$ -dimensional subproblems. At this point, we note that the space used by algorithm excepting the call to Thin-Projection is 2^n poly(\cdot). Therefore the total space used by the algorithm not including that used by Thin-Projection is 2^n poly(\cdot) as needed.

Given the above analysis, we see that the algorithm’s expected runtime (ignoring polynomial factors) is dominated by the recurrence relation

$$T(n) = (C_1g(n))^n + \max_{k \in [n]} (C_1g(n))^k + g(n)^k T(n - k)$$

where $T(0) = 1$ and $C_1 \geq 1$ is an absolute constant.

Claim: $T(n) \leq 2(n + 1)C_1^n g(n)^n$ Clearly the base case holds for $n = 0$. Next, since g is a non-decreasing, we have that

$$\begin{aligned} T(n) &= (C_1g(n))^n + \max_{k \in [n]} (C_1g(n))^k + g(n)^k T(n - k) \\ &\leq (C_1g(n))^n + \max_{k \in [n]} (C_1g(n))^k + g(n)^k 2(n - k + 1)(C_1g(n - k))^{n-k} \\ &\leq (C_1g(n))^n + (C_1g(n))^n + 2n(C_1g(n))^n = 2(n + 1)(C_1g(n))^n \end{aligned}$$

as needed. From the above relation, we have that the expected runtime of the algorithm is bounded by $2^{O(n)}g(n)^n \text{poly}(\cdot)$ as desired. \square

7.5 Convex Integer Minimization

In this section, we generalize our IP feasibility algorithms to solve convex integer minimization problems. More precisely, we give an algorithm to minimize any convex function f (admitting a subgradient oracle) over the lattice points inside a convex set. As noted in the introduction, if we are only interested in approximate minimizers then one can reduce the convex integer minimization problem to a sequence of feasibility problems via a standard binary search on the objective. With additional assumptions on the convex function f (linear, quasi-convex polynomial, etc.), such an approach often leads to an exact algorithm since we can bound the accuracy needed for binary search to find an optimal solution. In this section, we will show how to avoid all such “niceness” conditions on the function except convexity.

We note that binary search procedures require only the level sets of f (sets of the

form $f(\mathbf{x}) \leq M$) to be convex, i.e. they require only *quasi-convexity* of f . Indeed, our improvements here will be due to a fuller use of the convexity information of f .

We give the high level idea of the algorithm. To explain our approach, we overview the framework for our feasibility solvers and explain how to modify it to fit the optimization setting. Given a convex body K and lattice \mathcal{L} , our feasibility solver works as follows. First, we check if K is “flat” with respect to \mathcal{L} , i.e. either K has small lattice width or K admits a thin projection with respect to \mathcal{L} . If this is the case, we compute a decomposition of the integer points in K (via either hyperplane or subspace decompositions), and recurse on the lower dimensional subproblems. Notice that if we are minimizing a convex function f over $K \cap \mathcal{L}$, this step is still be valid, i.e. we compute the minimizers for all subproblems, and return the best one found. If K is not “flat”, we simply scale K until it becomes so (without comprimising the existence of an integer point in K), and now execute the recursion as in the previous step. This scaling step, on the other hand, is not valid in the optimization setting since it may remove the minimizing lattice point for f . Instead of scaling K until it becomes flat, in the optimization setting we will try to find a feasible solution $\mathbf{y} \in K \cap \mathcal{L}$ such that $K \cap \{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) \leq f(\mathbf{y})\}$ is flat. Since any solution improving on \mathbf{y} must lie in such a level set, it is valid to restrict the optimization problem in this way. Of course, the problem is how do we find a good feasible solution \mathbf{y} ? Here the idea is to solve use an approximate CVP solver to find a lattice point near the “center” of K (using algorithm Central-Lat-Pt from section 5.6). The main implication of K not being flat in this setting will be that the covering radius of K with respect to \mathcal{L} is $\ll 1$. Therefore K will be guaranteed to contain a “deep” lattice point \mathbf{y} . At this point, restricting to the level set of f induced by \mathbf{y} will not immediately make K flat. However, using the convexity of f , we will show that the restriction reduces the volume of K by a small factor (roughly $1 - 10^{-n}$). Using this volume decrease, we will be able to show that $2^{O(n)}$ iterations of this procedure will insure that K

indeed becomes flat. Once flatness has been achieved, we can use recursion as in the feasibility setting.

We now provide the exact algorithm below.

Algorithm 7.4 Convex-IP(K, f, \mathcal{L}, H)

Input: (\mathbf{a}_0, R) -circumscribed convex set $K \subseteq \mathbb{R}^n$ with separation oracle SEP_K , convex function $f : K \rightarrow \mathbb{R}$ equipped with subgradient oracle, n dimensional lattice \mathcal{L} with basis $B \in \mathbb{Q}^{n \times n}$, rational affine subspace $H = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$.

Output: Return NULL if $K \cap \mathcal{L} \cap H = \emptyset$ or $\mathbf{y} \in K \cap \mathcal{L} \cap H$ minimizing $f(\cdot)$.

- 1: $(K, \mathcal{L}, \mathbf{p}) \leftarrow \text{IP-Preprocess}(K, \mathcal{L}, H)$.
- 2: **if** $K = \emptyset$ **return** NULL; **else if** $\mathcal{L} = \{\mathbf{0}\}$ and $\mathbf{0} \in K$ **then return** \mathbf{p} .
- 3: $\mathbf{x} \leftarrow$ NULL.
- 4: **loop**
- 5: Build $(1 + \frac{1}{24k})$ -approx. HKZ basis $\mathbf{b}_1, \dots, \mathbf{b}_k$ for \mathcal{L} ($k = \dim(\mathcal{L})$) under $\|\cdot\|_{K-K}$.
- 6: **if** $\max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k}) \geq \frac{1}{3k}$ **then break** loop.
- 7: $(\mathbf{c}, \mathbf{y}) \leftarrow \text{Central-Lat-Pt}(K, \mathcal{L}, \frac{1}{3})$. $\mathbf{x} \leftarrow \mathbf{y} + \mathbf{p}$.
- 8: Retrieve $\mathbf{v} \in \partial f(\mathbf{x})$. **if** $\pi_{\text{span}(\mathcal{L})}(\mathbf{v}) = \mathbf{0}$ **then return** \mathbf{x} .
- 9: $K \leftarrow K \cap \{\mathbf{z} \in \mathbb{R}^n : f(\mathbf{z} + \mathbf{p}) \leq f(\mathbf{x})\}$.
- 10: $(K, \mathcal{L}, \mathbf{p}') \leftarrow \text{IP-Preprocess}(K, \mathcal{L}, \text{span}(\mathcal{L}))$. $\mathbf{p} \leftarrow \mathbf{p} + \mathbf{p}'$.
- 11: $j \leftarrow \arg \max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k})$. $W \leftarrow \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{j-1})^\perp \cap \text{span}(\mathcal{L})$.
- 12: **for all** $\mathbf{s} \in \text{Lattice-Enum}(\pi_W(K), \pi_W(\mathcal{L}), r)$ **do**
- 13: $\mathbf{y} \leftarrow \text{Convex-IP}(K, f(\cdot + \mathbf{p}), \mathcal{L}, \pi_W^{-1}(\mathbf{s}))$.
- 14: **if** $\mathbf{y} \neq \text{NULL}$ and $f(\mathbf{y} + \mathbf{p}) < f(\mathbf{x})$ **then** $\mathbf{x} \leftarrow \mathbf{y} + \mathbf{p}$.
- 15: **return** \mathbf{x} .

Theorem 7.5.1 (Correctness of Convex-IP). *On input K, f, \mathcal{L}, H as in 7.4, in expected $2^{O(n)}n^n \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space algorithm 7.4 either returns NULL if $K \cap \mathcal{L} \cap H = \emptyset$ or returns $\mathbf{y} \in \mathcal{L} \cap K \cap H$ minimizing $f(\cdot)$.*

We will require the following volumetric lemma for the analysis.

Lemma 7.5.2. *Let $K \subseteq \mathbb{R}^n$ be a n -dimensional γ -symmetric convex body. Then for any $\mathbf{x} \in K$, $\mathbf{v} \in \mathbb{R}^n$, letting $t = \langle \mathbf{v}, \mathbf{x} \rangle$, we have that*

$$\frac{\text{vol}_n(K \cap H_{\mathbf{v}, t}^{\geq})}{\text{vol}_n(K)} \geq \frac{1}{2} \gamma^n (1 - \|\mathbf{x}\|_K)^n$$

Proof. Let $K^+ = K \cap H_{\mathbf{v}, 0}^{\geq}$. First note that

$$\text{vol}_n(K^+) \geq \frac{1}{2} \text{vol}_n(K \cap -K) \geq \frac{1}{2} \gamma^n \text{vol}_n(K) \quad (7.5.1)$$

If $t \leq 0$ then $K^+ \subseteq K \cap H_{\mathbf{v},t}^{\geq}$, and hence the desired volume lower bound holds from equation (7.5.1).

Now assume that $t \geq 0$. Take $\mathbf{y} \in \arg \max_{\mathbf{z} \in K} \langle \mathbf{v}, \mathbf{z} \rangle$. Since $\frac{\mathbf{x}}{\|\mathbf{x}\|_K} \in K$, note that $0 \leq \frac{\langle \mathbf{v}, \mathbf{x} \rangle}{\|\mathbf{x}\|_K} \leq \langle \mathbf{v}, \mathbf{y} \rangle$.

I claim that $(1 - \|\mathbf{x}\|_K)K^+ + \|\mathbf{x}\|_K \mathbf{y} \subseteq K \cap H_{\mathbf{v},t}^{\geq}$. Take $\mathbf{w} = (1 - \|\mathbf{x}\|_K)\mathbf{z} + \|\mathbf{x}\|_K \mathbf{y}$, for $\mathbf{z} \in K^+$. Since $\mathbf{x} \in K \Rightarrow 0 \leq \|\mathbf{x}\|_K \leq 1$, we have that \mathbf{w} is a convex combination of points in K^+ , and hence $\mathbf{w} \in K^+$. Now note that

$$\langle \mathbf{v}, \mathbf{w} \rangle = (1 - \|\mathbf{x}\|_K) \langle \mathbf{v}, \mathbf{z} \rangle + \|\mathbf{x}\|_K \langle \mathbf{v}, \mathbf{y} \rangle \geq \|\mathbf{x}\|_K \langle \mathbf{v}, \mathbf{y} \rangle \geq \langle \mathbf{v}, \mathbf{x} \rangle = t.$$

Therefore we have that $\mathbf{w} \in H_{\mathbf{v},t}^{\geq}$ as needed..

By the Brunn-Minkowski inequality, we have that

$$\begin{aligned} \text{vol}_n(K \cap H_{\mathbf{v},t}^{\geq})^{\frac{1}{n}} &\geq \text{vol}_n((1 - \|\mathbf{x}\|_K)K^+ + \|\mathbf{x}\|_K \mathbf{y})^{\frac{1}{n}} \geq (1 - \|\mathbf{x}\|_K) \text{vol}_n(K^+)^{\frac{1}{n}} \\ &\geq 2^{-\frac{1}{n}} \gamma (1 - \|\mathbf{x}\|_K) \text{vol}_n(K)^{\frac{1}{n}} \end{aligned}$$

as needed. □

Analysis of Convex-IP.

We first discuss some implementation issues from the description of the algorithm.

Building an approximate HKZ basis: Here we proceed in the identical way as in Lemma 7.4.4. More precisely, to build a $(1 + \epsilon)$ -approximate HKZ basis for \mathcal{L} under $\|\cdot\|_{K-K}$, $\epsilon \in (0, 1)$, we proceed recursively by picking $\mathbf{b}_i^* \in \text{Shortest-Vector}(\pi_i(K - K), \pi_i(\mathcal{L}), \epsilon)$. Here we remember that π_i is the orthogonal projection onto $\text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{i-1})^\perp$.

Constructing a separation oracle for level sets of f : During the algorithm, we require separation oracles for bodies of the form $K' = K \cap \{\mathbf{z} \in \mathbb{R}^n : f(\mathbf{z}) \leq M\}$, where K is a convex subset of the domain of f .

To implement the separation oracle for K' we proceed as follows. When queried on a point $\mathbf{x} \in \mathbb{Q}^n$, we first check that $\mathbf{x} \in K$ and then check that $f(\mathbf{x}) \leq M$. If $\mathbf{x} \notin K$, then we use the separation oracle for K to separate \mathbf{x} from K' . If $\mathbf{x} \in K$ and $f(\mathbf{x}) > M$, we use a subgradient $\mathbf{v} \in \partial f(\mathbf{x})$ (which we have query access to) as our separator. Here we note that by subgradient properties of \mathbf{v} , that $\langle \mathbf{v}, \mathbf{z} \rangle \geq \langle \mathbf{v}, \mathbf{x} \rangle$ implies that $f(\mathbf{z}) \geq f(\mathbf{x}) > M$, and hence $\mathbf{z} \notin K'$. Therefore $\sup_{\mathbf{z} \in K'} \langle \mathbf{v}, \mathbf{z} \rangle < \langle \mathbf{v}, \mathbf{x} \rangle$ as needed.

Correctness: For the first step, we run our standard preprocessing algorithm on K, \mathcal{L}, H to reduce to a full dimensional problem. Here we obtain a new K and \mathcal{L} such that $(K \cap \mathcal{L}) + \mathbf{p}$ corresponds exactly to the original solutions, $\dim(K) = \dim(\mathcal{L})$, and K is well sandwiched. If the preprocessing returns that $K = \emptyset$, then the set of solutions is empty, and hence we correctly return NULL. If the preprocessing returns $\mathcal{L} = \{\mathbf{0}\}$ and $\mathbf{0} \in K$, then the set of solutions corresponds to a single point, i.e. \mathbf{p} , and so we correctly return \mathbf{p} as a minimizer.

Next, we initialize $\mathbf{x} \leftarrow \text{NULL}$, which will denote the best solution found thus far. We now begin the first main loop. We shall establish the following invariants for this loop:

- (1) At the beginning of each iteration, any improving solution over \mathbf{x} must be contained in $(K \cap \mathcal{L}) + \mathbf{p}$. Furthermore, every solution $\mathbf{z} \in (K \cap \mathcal{L}) + \mathbf{p}$ satisfies $f(\mathbf{z}) \leq f(\mathbf{x})$ (with the convention $f(\mathbf{x}) = \infty$ if $\mathbf{x} = \text{NULL}$).
- (2) At the end of any non-terminal iteration, either the volume of K drops by at least a $(1 - \frac{1}{2}10^{-\dim(K)})$ factor or the dimension of K decreases.
- (3) If the loop terminates, then either we have found a valid minimizer for f or we have found an l -dimensional subspace W s.t. $G(\pi_W(K), \pi_W(\mathcal{L}))^{\frac{1}{l}} \leq 7 \dim(\mathcal{L})$.

Proof of Loop Invariants. At the start of the first iteration, invariant (1) clearly holds

since $(K \cap \mathcal{L}) + \mathbf{p}$ corresponds to the entire solution space and $\mathbf{x} = \text{NULL}$. We now assume that the invariants hold at the beginning of the current loop iteration, and show that they are maintained at the beginning of the next iteration.

At the start of a loop iteration, we build an $(1 + \frac{1}{24k})$ -approximate HKZ basis b_1, \dots, b_k for \mathcal{L} (as described above). We note that running Shortest-Vectors on $K - K$ and \mathcal{L} is possible here since we have explicit sandwiching guarantees for K in $\text{span}(\mathcal{L})$ (retrieved from IP-Preprocess).

From here if $\max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k}) \geq \frac{1}{3k}$, we break from the loop. In this case, upon exiting the loop, we set $j \leftarrow \arg \max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k})$, and let $W = \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_{j-1})^\perp \cap \text{span}(\mathcal{L})$. Letting $l = k - j + 1$, we see that W is a l -dimensional subspace of \mathbb{R}^n . We claim that $G(\pi_W(K), \pi_W(\mathcal{L}))^{\frac{1}{l}} \leq 7k \leq 7 \dim(\mathcal{L})$. To see this, note that

$$\frac{1}{3k} \leq D_{\pi_W(K-K)}(\mathbf{b}_j^*, \frac{1}{24k}) \leq (1 + \frac{1}{24k}) \|\mathbf{b}_j^*\|_{\pi_W(K-K)} \leq (1 + \frac{1}{24k})^2 \lambda_1(\pi_W(K-K), \pi_W(\mathcal{L}))$$

For $k \geq 1$, the above implies that $\lambda_1(\pi_W(K-K), \pi_W(\mathcal{L})) \geq \frac{3}{10k}$. As in the proof of Theorem 7.4.3, by Minkowski's first theorem this implies that

$$\frac{\text{vol}_l(\pi_W(K))^{\frac{1}{l}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{l}}} \leq \frac{10k}{3}$$

Assume first that $\mu(K, \mathcal{L}) \leq 1$. Then by Corollary 7.4.2 we have that

$$G(\pi_W(K), \pi_W(\mathcal{L}))^{\frac{1}{l}} \leq 2 \frac{\text{vol}_l(\pi_W(K))^{\frac{1}{l}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{l}}} \leq 2 \frac{10k}{3} \leq 7k$$

as needed. Assume that $\mu(K, \mathcal{L}) \geq 1$. Letting $\mu = \mu(K, \mathcal{L})$, we have that

$$G(\pi_W(K), \pi_W(\mathcal{L})) \leq G(\pi_W(\mu K), \pi_W(\mathcal{L})) \quad (\text{since } \mu \geq 1),$$

and hence it suffices to upper bound the latter term. Here we note that $\mu(\mu K, \mathcal{L}) = \frac{\mu(K, \mathcal{L})}{\mu} = 1$. Now using the exact same argument as in Lemma 7.4.4, remembering that $\pi_j = \pi_W$, we have that

$$1 = \mu(\mu K, \mathcal{L}) \leq \sum_{i=1}^k \|\mathbf{b}_i^*\|_{\pi_i(\mu(K-K))} \leq k \frac{24k+1}{24k-1} \|\mathbf{b}_j^*\|_{\pi_W(\mu(K-K))}$$

Using the fact that $\|\mathbf{b}_j^*\|_{\pi_W(\mu(K-K))} \leq (1 + \frac{1}{24k})\lambda_1(\pi_W(\mu(K-K)), \pi_W(\mathcal{L}))$, using the same argument as in Lemma 7.4.4, we get that

$$\left(\sum_{i=1}^k \|\mathbf{b}_i^*\|_{\pi_i(\mu(K-K))} \right) \frac{\text{vol}_l(\pi_W(\mu K))^{\frac{1}{l}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{l}}} \leq k \left(\frac{24k+1}{24k} \right) \left(\frac{24k+1}{24k-1} \right)$$

Since $\sum_{i=1}^k \|\mathbf{b}_i^*\|_{\pi_i(\mu(K-K))} \geq 1$, the above inequality yields

$$\frac{\text{vol}_l(\pi_W(\mu K))^{\frac{1}{l}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{l}}} \leq k \left(\frac{24k+1}{24k} \right) \left(\frac{24k+1}{24k-1} \right) \leq \frac{8}{7}k$$

for $k \geq 1$. Since $\mu(\mu K, \mathcal{L}) = 1$ by Corollary 7.4.2 we have that

$$G(\pi_W(K), \pi_W(\mathcal{L}))^{\frac{1}{l}} \leq 2 \frac{\text{vol}_l(\pi_W(K))^{\frac{1}{l}}}{\det(\pi_W(\mathcal{L}))^{\frac{1}{l}}} \leq 2 \frac{8k}{7} \leq 7k$$

as needed. Therefore, if we the loop terminates because $\max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k}) \geq \frac{1}{3k}$, then we have indeed found a “good” projection subspace W as needed. This takes care of the second part of loop invariant (3).

If instead $\max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k}) < \frac{1}{3k}$, we get that $\max_{i \in [k]} \|\mathbf{b}_i^*\|_{\pi_i(K-K)} < \frac{1}{3k} + \frac{1}{24k} = \frac{9}{24k}$. From Lemma 2.4.7, we know that

$$\mu(K, \mathcal{L}) \leq \sum_{i=1}^k \|\mathbf{b}_i^*\|_{\pi_i(K-K)} < k \frac{9}{24k} = \frac{9}{24}$$

Let $(\mathbf{c}, \mathbf{y}) \leftarrow \text{Central-Lat-Pt}(K, \mathcal{L}, \frac{1}{3})$. By the guarantees on algorithm Central-Lat-Pt (see Theorem 5.6.1), we have that $K - \mathbf{c}$ is $\frac{1}{5}$ -symmetric and \mathbf{y} satisfies

$$\|\mathbf{y} - \mathbf{c}\|_{K-\mathbf{c}} \leq (1 + \frac{1}{3})d_{K-\mathbf{c}}(\mathcal{L}, \mathbf{c}) \leq \frac{4}{3}\mu(K, \mathcal{L}) = \frac{1}{2} \quad (7.5.2)$$

From the above, we in particular have that $\mathbf{y} \in K \cap \mathcal{L}$. By loop invariant (1), we immediately have that $f(\mathbf{y} + \mathbf{p}) \leq f(\mathbf{x})$. Therefore, setting $\mathbf{x} \leftarrow \mathbf{y} + \mathbf{p}$, \mathbf{x} becomes the best solution found thus far, and by the furthermore of loop invariant (1), any improving solution must be contained $(K \cap \mathcal{L}) + \mathbf{p}$.

Now we take $\mathbf{v} \in \partial f(\mathbf{x})$. If $\pi_{\text{span}(\mathcal{L})}(\mathbf{v}) = \mathbf{0}$, we claim that \mathbf{x} is a global minimizer for f restricted to the affine subspace $\mathbf{x} + \text{span}(\mathcal{L})$ (and hence in particular on $(K \cap \mathcal{L}) + \mathbf{p}$).

To see this, take $\mathbf{z} \in \mathbf{x} + \text{span}(\mathcal{L})$. Then since \mathbf{v} is a subgradient, we have that $f(\mathbf{z}) \geq f(\mathbf{x}) + \langle \mathbf{v}, \mathbf{z} - \mathbf{x} \rangle$. Since $\mathbf{z} - \mathbf{x} \in \text{span}(\mathcal{L})$, we have that $\langle \mathbf{v}, \mathbf{z} - \mathbf{x} \rangle = 0$ and hence $f(\mathbf{z}) \geq f(\mathbf{x})$, as needed. From here we see that \mathbf{x} is at least as good any potential improving solution (again by loop invariant (1)), and hence we correctly return \mathbf{x} as an optimal solution for the IP. This completes the proof of invariant (3).

Now assume that $\pi_{\text{span}(\mathcal{L})}(\mathbf{v}) \neq \mathbf{0}$. Let $K' \leftarrow K \cap \{\mathbf{z} \in \mathbb{R}^n : f(\mathbf{z} + \mathbf{p}) \leq f(\mathbf{x})\}$. Since any improving solution over \mathbf{x} must be in $(K \cap \mathcal{L}) + \mathbf{p}$ and have better objective value than \mathbf{x} , we clearly have that any improving solution must also be in $(K' \cap \mathcal{L}) + \mathbf{p}$. Letting $t = \langle \mathbf{v}, \mathbf{y} \rangle$, we note that

$$H_{\mathbf{v},t}^{\leq} = \{\mathbf{z} \in \text{span}(\mathcal{L}) : \langle \mathbf{v}, \mathbf{z} \rangle \leq t\}$$

is a non-trivial supporting halfspace for K' . Now by construction $\mathbf{y} \in H_{\mathbf{v},t}$ and $\dim(K) = \dim(\mathcal{L}) = k$. Since $K - \mathbf{c}$ is $\frac{1}{5}$ -symmetric and $\|\mathbf{y} - \mathbf{c}\|_{K-\mathbf{c}} \leq \frac{1}{2}$ by Lemma 7.5.2 we have that

$$\begin{aligned} \text{vol}_k(K') &\leq \text{vol}_k(K \cap H_{\mathbf{v},t}^{\leq}) = \text{vol}_k(K) - \text{vol}_k(K \setminus H_{\mathbf{v},t}^{\leq}) \\ &\leq \text{vol}_k(K) - \frac{1}{2} 5^{-k} 2^{-k} \text{vol}_k(K) = (1 - \frac{1}{2} 10^{-k}) \text{vol}_k(K) \end{aligned} \tag{7.5.3}$$

From the above, we see that letting $K \leftarrow K'$ (as is done on line 9) decreases the volume of K by at least a $(1 - \frac{1}{2} 10^{-\dim(K)})$ factor. Next, we run our preprocessing algorithm on K , \mathcal{L} , and $\text{span}(\mathcal{L})$ to recover updated $(K, \mathcal{L}, \mathbf{p}')$. We remember that the set of solutions $(K \cap \mathcal{L}) + \mathbf{p}'$ after preprocessing corresponds exactly to $(K \cap \mathcal{L})$ before preprocessing. Therefore the update $\mathbf{p} \leftarrow \mathbf{p} + \mathbf{p}'$, leaves the set $(K \cap \mathcal{L}) + \mathbf{p}$ unchanged before and after preprocessing, and hence loop invariant (1) is still satisfied (and remains satisfied at the beginning of the next iteration). If $\dim(K)$ decreases after preprocessing, then loop invariant (2) is satisfied. Since $K + \mathbf{p}'$ after preprocessing is a subset of K before preprocessing, if $\dim(K)$ is unchanged by the preprocessing, then as argued above the volume of K has indeed dropped as required. Therefore loop invariant (2) is satisfied. This completes the proof of the loop invariants. \square

After the first main loop, the remainder of the algorithm is straightforward. Here we decompose the remainder of the (shifted) solution space (which contains any improving solution over \mathbf{x}) using the projection onto W . More precisely we use the decomposition $K \cap \mathcal{L} \subseteq \bigcup_{\mathbf{s} \in \pi_W(K) \cap \pi_W(\mathcal{L})} \pi_W^{-1}(\mathbf{s}) \cap \mathcal{L}$, using Lattice-Enum to iterate over the set $\pi_W(\mathcal{L}) \cap \pi_W(K)$. During the iteration, for each $\mathbf{s} \in \pi_W(K) \cap \pi_W(\mathcal{L})$, we recursively solve the lower dimensional convex integer minimization problem associated with K, \mathcal{L} and $\pi_W^{-1}(\mathbf{s})$ and $f(\cdot + \mathbf{p})$. As we iterate, we keep track of the best solution found over all recursive calls (shifting by \mathbf{p} each time), and return this solution at the end of the loop. This completes the proof of correctness.

Runtime: The initial preprocessing step on K, \mathcal{L}, H takes $2^{O(n)} \text{poly}(\cdot)$ time and $2^n \text{poly}(\cdot)$ space.

Next, we need to bound the time spent in the first main loop. During each loop iteration, in the worst case, we build a new approximate HKZ basis for \mathcal{L} with respect to $\|\cdot\|_{K-K}$, make a call to Central-Lat-Pt on K and \mathcal{L} , and run the standard preprocessing on a restriction of K and \mathcal{L} . By the guarantees on each invoked subroutine, the execution of all these steps takes at most $2^{O(\dim(\mathcal{L}))} \text{poly}(\cdot) = 2^{O(n)} \text{poly}(\cdot)$ expected time and $2^{\dim(\mathcal{L})} \text{poly}(\cdot) = 2^{O(n)} \text{poly}(\cdot)$ space (notice that the randomized runtime comes only from Central-Lat-Pt).

We now show that the maximum number of iterations of the first main loop is $2^{O(n)} \text{poly}(\cdot)$. To do this, we give a bound on maximum number of consecutive iterations during which the volume of K can be decreased by a $(1 - \frac{1}{2}10^{-k})$ factor, for $k = \dim(K)$, without either finding a “good” projection, decreasing the dimension of K , or finding a valid minimizer for f . At the start of any such sequence of iterations, we note that K is (\mathbf{a}'_0, R') -circumscribed, for some $R' \leq R$ (where R is the radius of the original circumscribing ball). Therefore at the beginning of this sequence, $\text{vol}_k(K) \leq \text{vol}_k(R'B_2^k) \leq (2R)^k$. Then after T iterations, letting $\lambda = \frac{1}{2}10^{-k}$, we have

that

$$\text{vol}_k(K) \leq (1 - \frac{1}{2}10^{-k})^T (2R)^k \leq e^{-\lambda T} (2R)^k.$$

Then after $T = \max\{\lceil \frac{\ln((2R)^k / \det(\mathcal{L}))}{\lambda} \rceil + 1, 1\}$ iterations, we have that

$$\text{vol}_k(K) \leq e^{-\lambda T} (2R)^k < \det(\mathcal{L})$$

We note that since $\mathcal{L} \subseteq \mathbb{Q}^n$, and \mathcal{L} is a sublattice of the original input lattice, we have that $|\ln(\det(\mathcal{L}))|$ is polynomial in the size of the input basis. Therefore $T = \text{poly}(\cdot) \frac{1}{\lambda} = \text{poly}(\cdot) 10^k = \text{poly}(\cdot) 10^n$. We claim that no such sequence can last more than $10^n \text{poly}(\cdot)$ iterations. To prove this, it suffices to show that if $\text{vol}_k(K) < \det(\mathcal{L})$ at the beginning of a loop iteration, then we guaranteed to find a “good” projection for K . To begin, by Lemma 2.4.4 we know that if K is \mathcal{L} -covering (i.e. $\mu(K, \mathcal{L}) \leq 1$), then $\text{vol}_k(K) \geq \det(\mathcal{L})$. Since $\text{vol}_k(K) < \det(\mathcal{L})$, we must have that K is not \mathcal{L} -covering, and hence $\mu(K, \mathcal{L}) \geq 1$. Now let $\mathbf{b}_1, \dots, \mathbf{b}_k$ denote the $(1 + \frac{1}{24k})$ -approximate HKZ basis computed during the loop iteration. Then we see that

$$1 \leq \mu(K, \mathcal{L}) \leq \sum_{i=1}^k \|\mathbf{b}_i^*\|_{\pi_i(K-K)} \leq k \max_{i \in [k]} \|\mathbf{b}_i^*\|_{\pi_i(K-K)}.$$

Hence $\frac{1}{k} \geq \max_{i \in [k]} \|\mathbf{b}_i^*\|_{\pi_i(K-K)}$. Therefore by the guarantees on the distance oracles, we have that $\max_{i \in [k]} D_{\pi_i(K-K)}(\mathbf{b}_i^*, \frac{1}{24k}) \geq \frac{1}{k} - \frac{1}{24k} \geq \frac{1}{3k}$. Given this, we exit the first main loop upon the check at line 6, and hence we find good projection for K as desired. From here, we see that during any sequence of $10^n \text{poly}(\cdot)$ iterations, we must either find a good projection, decrease the dimension of K , or find a valid minimizer for f . Note that if we find a good projection, or find a valid minimizer, we exit the loop. On the other hand if we decrease the dimension of K (which happens during IP-Preprocess), we decrease dimension by at least 1. Since at the beginning of the first main loop, the dimension of K is at most n , we cannot decrease dimension more than n times. Therefore, a bound on the maximum number of iterations of the first main loop is $n 10^n \text{poly}(\cdot) = 2^{O(n)} \text{poly}(\cdot)$, as needed. Therefore to finish executing

the entire first main loop requires at most $2^{O(n)}$ poly(\cdot) expected time and 2^n poly(\cdot) space.

After the first main loop, assuming we do not find a valid minimizer for f , we need to execute all the recursive subproblems indexed by $\pi_W(K) \cap \pi_W(\mathcal{L})$. By loop invariant (2), we note that for the computed projection subspace W , we have that $G(\pi_W(K), \pi_W(\mathcal{L})) \leq (7 \dim(\mathcal{L}))^{\dim(W)} \leq (7n)^{\dim(W)}$. At this point, the recursion relation for the runtime is essentially identical to that of Algorithm 7.2 (with expected running time replacing deterministic running time). Therefore using the identical analysis as Theorem 7.1.2, we get that entire algorithm runs in $2^{O(n)}n^n$ poly(\cdot) expected time using 2^n poly(\cdot) space, as needed. \square

7.6 Conclusion

The Integer Programming Problem is one of the central computational problems in computer science and operations research. The study of algorithms for IP has been an active area of research since the 1950's, and a plethora of algorithmic techniques have been developed to solve IP from both a practical and theoretical standpoint. In this Chapter, we have developed new geometric techniques for solving integer programs, based on solvers for general norm lattice problems. With these methods, we have implemented new Lenstra and Kannan type algorithms, which solve n dimensional integer programs in $2^{O(n)}(n^{\frac{4}{3}} \text{polylog}(n))^n$ and $2^{O(n)}n^n$ time using $O(2^n)$ space respectively. These algorithms represent the fastest known algorithms of either class, and in particular the fastest known theoretically efficient algorithms for IP. Furthermore, our algorithms work for any integer program whose feasible region is specified by separation oracle, even when the feasible region is not full dimensional, generalizing over previous approaches.

As mentioned in introduction, a major open question is whether there exists a $2^{O(n)}$ time algorithm for IP. Though our algorithms retain the same asymptotic $n^{O(n)}$

running times of previous algorithms, we believe they represent significant milestones towards finding a single exponential time algorithm.

Firstly, our use (and development) of general norm solvers for lattice problems, has allowed us to avoid the complexity blowups due to ellipsoidal approximations used by previous algorithms. These approximations almost inevitably lead to at least $n^{\frac{n}{2}}$ increases in running time, and hence represented a first major bottleneck to be overcome in search of a $2^{O(n)}$ algorithm.

Having surmounted this obstacle, our Kannan type algorithm allowed us to make headway on another bottleneck, i.e. the reliance on Kinchine’s flatness theorem. In the worst case, it is well-known that a lattice free convex body may have lattice width n . Therefore, the number of subproblems generated by any Lenstra type algorithm on such a body must be $\geq n$. Though one does not expect such worst case behavior at every level of recursion, it nevertheless seems unlikely that a hyperplane branching scheme can achieve better than $n^{\Omega(n)}$ running time. Our Kannan type algorithm however, is based on a completely different transference theorem of Kannan and Lovasz [67] (see Theorem 7.4.3), which does not suffer from this type of lower bound. In fact, the only nontrivial lower bound for the quality of a “thin projection” is $\Omega(\log n)$. This bound was conjectured to be tight by Kannan and Lovász [67] (see Conjecture 7.4.5), and we show in Theorem 7.3 that a suitable algorithmic version of this conjecture would yield an $O(\log n)^n$ time algorithm for IP.

Future Research. As a first future research direction, we wish to improve the current bounds on the quality of thin projections. The current best bound of n , due to Kannan and Lovasz, is based on the use of generalized HKZ bases. However, even when $K = B_2^n$, there are examples of HKZ bases which yield projections of quality $O(\sqrt{n})$. Therefore it seems that substantially new techniques are required in this setting. One potentially promising direction in this effort is to see whether the

discrete gaussian technology developed for proving Kinchine's flatness theorem can be brought to bear on this problem.

Another research direction, pertains to getting a better analysis of either our Kannan or Lenstra type algorithm. In particular, we currently assume the worst case bounds at every recursion node to obtain our complexity bounds. A main question here is whether this is really possible, and in particular if a more careful analysis could allow us to improve the complexity of the algorithms. On the flip side, as mentioned above, it seems unlikely that any Lenstra type algorithm can in general solve IP in faster than $n^{\Omega(n)}$ time. Here it would be very interesting to show that any algorithm which proceeds only by branching on hyperplanes must take at least $n^{\Omega(n)}$ time.

REFERENCES

- [1] AJTAI, M., “The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract),” in *STOC*, pp. 10–19, 1998.
- [2] AJTAI, M., KUMAR, R., and SIVAKUMAR, D., “A sieve algorithm for the shortest lattice vector problem,” in *STOC*, pp. 601–610, 2001.
- [3] AJTAI, M., KUMAR, R., and SIVAKUMAR, D., “Sampling short lattice vectors and the closest lattice vector problem,” in *IEEE Conference on Computational Complexity*, pp. 53–57, 2002.
- [4] ARORA, S., BABAI, L., STERN, J., and SWEEDYK, Z., “The hardness of approximate optima in lattices, codes, and systems of linear equations,” *J. Comput. Syst. Sci.*, vol. 54, no. 2, pp. 317–331, 1997. Preliminary version in FOCS 1993.
- [5] ARVIND, V. and JOGLEKAR, P. S., “Some sieving algorithms for lattice problems,” in *FSTTCS*, pp. 25–36, 2008.
- [6] AVIS, D. and FUKUDA, K., “Reverse search for enumeration,” *Discrete Applied Mathematics*, vol. 65, pp. 21–46, 1993.
- [7] BABAI, L., “On Lovász’ lattice reduction and the nearest lattice point problem,” *Combinatorica*, vol. 6, no. 1, pp. 1–13, 1986. Preliminary version in STACS 1985.
- [8] BALL, K. M., “Logarithmically concave functions and sections of convex sets in r^n ,” *Studia Mathematica*, vol. 88, pp. 69–84, 1988.
- [9] BALL, K. M., “An elementary introduction to modern convex geometry,” in *Flavor of Geometry, Number 31 in MSRI Publications* (LEVY, S., ed.), pp. 1–58, Cambridge University Press, 1997.
- [10] BANASZCZYK, W., “New bounds in some transference theorems in the geometry of numbers,” *Mathematische Annalen*, vol. 296, pp. 625–635, 1993.
- [11] BANASZCZYK, W., “Inequalities for convex bodies and polar reciprocal lattices in R^n ,” *Discrete and Computational Geometry*, vol. 13, pp. 217–231, 1995.
- [12] BANASZCZYK, W., “Inequalities for convex bodies and polar reciprocal lattices in R^n II: Application of k -convexity,” *Discrete and Computational Geometry*, vol. 16, pp. 305–311, 1996.

- [13] BANASZCZYK, W., LITVAK, A., PAJOR, A., and SZAREK, S., “The flatness theorem for nonsymmetric convex bodies via the local theory of banach spaces,” *Mathematics of Operations Research*, vol. 24, no. 3, pp. 728–750, 1999.
- [14] BARVINOK, A., “A polynomial time algorithm for counting integral points in polyhedra when the dimension is fixed,” *Mathematics of Operations Research*, vol. 19, no. 4, pp. 769–779, 1994.
- [15] BEN-TAL, A. and NEMIROVSKI, A., *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2001.
- [16] BILLINGSLEY, P., *Probability and Measure*. Wiley-Interscience, 3rd ed. ed., 1995.
- [17] BLASCHKE, W., “Über affine geometry xiv: eine minimum aufgabe für legendres trägheits ellipsoid,” *Ber. verh. sächs. Akad. d. Wiss.*, vol. 70, pp. 72–75, 1918.
- [18] BLÖMER, J. and NAEWE, S., “Sampling methods for shortest vectors, closest vectors and successive minima,” in *ICALP*, pp. 65–77, 2007.
- [19] BONAMI, P., DASH, G. C. S., FISCHETTI, M., and LODI, A., “Projected Chvatal-Gomory Cuts for Mixed Integer Linear Programs,” *Mathematical Programming*, vol. 113, pp. 241–257, 2008.
- [20] BOURGAIN, J., “On high-dimensional maximal functions associated to convex bodies,” *Amer. J. Math.*, vol. 108, no. 6, pp. 1467–1476, 1986.
- [21] BOURGAIN, J. and MILMAN, V. D., “New volume ratio properties for convex symmetric bodies in \mathbb{R}^n ,” *Inventiones Mathematicae*, vol. 88, pp. 319–340, 1987.
- [22] BROWDER, A., *Mathematical Analysis: An Introduction*. Springer-Verlag, 1995.
- [23] CAI, J.-Y. and NERURKAR, A., “Approximating the SVP to within a factor $(1+1/\dim^c)$ is NP-hard under randomized reductions,” *J. Comput. Syst. Sci.*, vol. 59, no. 2, pp. 221–239, 1999. Preliminary version in CCC 1998.
- [24] CASSELS, J. W. S., *An introduction to Diophantine approximation*. New York: Hafner, 1972.
- [25] CASSELS, J., *An Introduction to the Geometry of Numbers*. Springer Verlag, 1971.
- [26] ÇEZİK, M. T. and IYENGAR, G., “Cuts for mixed 0-1 conic programming,” *Mathematical Programming*, vol. 104, pp. 179–202, 2005.
- [27] CHUNG, K., DADUSH, D., LIU, F., and PEIKERT, C., “On the smoothing parameter problem of a lattice.” Preprint, 2012.

- [28] CHVÁTAL, V., “Edmonds polytopes and a hierarchy of combinatorial problems,” *Discrete Mathematics*, vol. 4, pp. 305–337, 1973.
- [29] CORMEN, T., LEISERSON, C., RIVEST, R., and STEIN, C., *Introduction to Algorithms*. The MIT Press, 3rd ed. ed., 2009.
- [30] DADUSH, D., DEY, S. S., and VIELMA, J. P., “The Chvátal-Gomory Closure of Strictly Convex Body.” To appear in *Mathematics of Operations Research*, 2010.
- [31] DADUSH, D. and VEMPALA, S., “Deterministic construction of an approximate m-ellipsoid and its application to derandomizing lattice algorithms,” in *SODA*, 2012.
- [32] DADUSH, D. and VEMPALA, S., “Near-optimal deterministic algorithms for volume computation and lattice problems via m-ellipsoids.” Preprint, 2012.
- [33] DADUSH, D., “A $o(\frac{1}{\epsilon^2})$ -time algorithm for approximate integer programming,” in *LATIN 2012*, 2012.
- [34] DADUSH, D., DEY, S. S., and VIELMA, J. P., “On the chvátal-gomory closure of a compact convex set,” in *Integer Programming and Combinatorial Optimization (IPCO)*, pp. 130–142, 2011.
- [35] DADUSH, D., DEY, S. S., and VIELMA, J. P., “On the split closure of a strictly convex body,” *Operations Research Letters*, vol. 39, no. 2, pp. 227–239, 2011.
- [36] DADUSH, D., PEIKERT, C., and VEMPALA, S., “Enumerative lattice algorithms in any norm via m-ellipsoid coverings,” in *FOCS*, 2011.
- [37] DANTZIG, G., FULKERSON, R., and JOHNSON, S., “Solution of a large scale traveling salesmen problem,” *Operations Research*, vol. 2, pp. 393–410, 1954.
- [38] DEY, S. S. and VIELMA, J. P., “The Chvátal-Gomory Closure of an Ellipsoid Is a Polyhedron,” in *IPCO XIV*, pp. 327–340, 2010.
- [39] DIESTEL, R., *Graph Theory*, vol. 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, 4th ed. ed., 2010.
- [40] DINUR, I., “Approximating SVP_∞ to within almost-polynomial factors is NP-hard,” *Theor. Comput. Sci.*, vol. 285, no. 1, pp. 55–71, 2002. Preliminary version in CIAC 2000.
- [41] DINUR, I., KINDLER, G., RAZ, R., and SAFRA, S., “Approximating CVP to within almost-polynomial factors is NP-hard,” *Combinatorica*, vol. 23, no. 2, pp. 205–243, 2003. Preliminary version in FOCS 1998.
- [42] D.S. DUMMIT, R. F., *Abstract Algebra*. Hoboken, New Jersey: John Wiley & Son, 3rd ed. ed., 2004.

- [43] DUNKEL, J. and SCHULZ, A. S., “The Gomory-chvátal closure of a non-rational polytope is a rational polytope.” http://www.optimization-online.org/DB_HTML/2010/11/2803.html, 2010.
- [44] DYER, M. E., FRIEZE, A. M., and KANNAN, R., “A random polynomial time algorithm for approximating the volume of convex bodies,” in *STOC*, pp. 375–381, 1989.
- [45] DYER, M., FRIEZE, A., and KANNAN, R., “A random polynomial-time algorithm for approximating the volume of convex bodies,” *J. ACM*, vol. 38, no. 1, pp. 1–17, 1991.
- [46] EDMONDS, J., “Paths, trees, and flowers,” *Canadian Journal of mathematics*, vol. 17, pp. 449–467, 1965.
- [47] EISENBRAND, F. and SHMONIN, G., “Parametric integer programming in fixed dimension,” *Mathematics of Operations Research*, vol. 33, no. 4, pp. 839–850, 2008.
- [48] ELEKES, G., “A geometric inequality and the complexity of computing volume,” *Discrete & Computational Geometry*, pp. 289–292, 1986.
- [49] FIGIEL, T. and TOMCZAK-JAEGERMANN, N., “Projections onto hilbertian subspaces of banach spaces,” *Israel Journal of Mathematics*, vol. 33, pp. 155–171, 1979.
- [50] FISCHETTI, M. and LODI, A., “Optimizing over the first Chvátal closure,” *Mathematical Programming, Series B*, vol. 110, pp. 3–20, 2007.
- [51] FUREDI, Z. and BARANY, I., “Computing the volume is difficult,” in *STOC ’86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*, (New York, NY, USA), pp. 442–447, ACM, 1986.
- [52] FUREDI, Z. and BARANY, I., “Approximation of the sphere by polytopes having few vertices,” *Proceedings of the AMS*, vol. 102, no. 3, 1988.
- [53] GIANNOPOULOS, A. A. and MILMAN, V. D., “Chapter 17: Euclidean structure in finite dimensional normed spaces,” in *Handbook of the Geometry of Banach Spaces* (JOHNSON, W. and LINDENSTRAUSS, J., eds.), vol. 1, pp. 707–779, Elsevier Science B.V., 2001.
- [54] GILES, J., *Introduction to the Analysis of Normed Linear Spaces*. Cambridge University Press, 2000.
- [55] GOMORY, R., “An outline of an algorithm for solving integer programs,” *Bulletin of the American Mathematical Society*, vol. 64, no. 5, pp. 275–278, 1958.
- [56] GRÖTSCHEL, M., LOVÁSZ, L., and SCHRIJVER, A., *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1988.

- [57] GRÖTSCHEL, M. and PADBERG, M., “On the symmetric travelling salesman problem I: Inequalities,” *Math. Programming*, vol. 16, pp. 265–280, 1979.
- [58] GRÖTSCHEL, M. and PADBERG, M., “On the symmetric travelling salesman problem II: Lifting theorems and facets,” *Math. Programming*, vol. 16, pp. 281–302, 1979.
- [59] GRUBER, P. and LEKKERKERKER, C., *Geometry of Numbers*. Amsterdam - New York: North Holland, 2nd ed. ed., 1987.
- [60] HALMOS, P., *Measure Theory*, vol. 18 of *Graduate Texts in Mathematics*. New York: Springer-Verlag, 1950.
- [61] HANROT, G. and STEHLÉ, D., “Improved analysis of kannan’s shortest lattice vector algorithm,” in *Proceedings of the 27th annual international cryptology conference on Advances in cryptology, CRYPTO’07*, (Berlin, Heidelberg), pp. 170–186, Springer-Verlag, 2007.
- [62] HAVIV, I. and REGEV, O., “Tensor-based hardness of the shortest vector problem to within almost polynomial factors,” in *STOC*, pp. 469–477, 2007.
- [63] HEINZ, S., “Complexity of integer quasiconvex polynomial optimization,” *Journal of Complexity*, vol. 21, no. 4, pp. 543 – 556, 2005. Festschrift for the 70th Birthday of Arnold Schonhage.
- [64] HELFRICH, B., “Algorithms to construct minkowski reduced and hermite reduced lattice bases,” *Theor. Comput. Sci.*, vol. 41, pp. 125–139, December 1985.
- [65] HILDEBRAND, R. and KÖPPE, M., “A new lenstra-type algorithm for quasiconvex polynomial integer minimization with complexity $2^{O(n \log n)}$.” Arxiv, Report 1006.4661, 2010. <http://arxiv.org>.
- [66] JOUX, A. and STERN, J., “Lattice reduction: A toolbox for the cryptanalyst,” *J. Cryptology*, vol. 11, no. 3, pp. 161–185, 1998.
- [67] KANNAN, R. and LOVÁSZ, L., “Covering minima and lattice point free convex bodies,” *Annals of Mathematics*, vol. 128, pp. 577–602, 1988.
- [68] KANNAN, R., LOVÁSZ, L., and SIMONOVITS, M., “Isoperimetric problems for convex bodies and a localization lemma,” *Discrete & Computational Geometry*, vol. 13, pp. 541–559, 1995.
- [69] KANNAN, R., LOVÁSZ, L., and SIMONOVITS, M., “Random walks and an $O^*(n^5)$ volume algorithm for convex bodies,” *Random Structures and Algorithms*, vol. 11, pp. 1–50, 1997.
- [70] KANNAN, R., “Minkowski’s convex body theorem and integer programming,” *Mathematics of operations research*, vol. 12, pp. 415–440, August 1987. 1987.

- [71] KANNAN, R., “Test sets for integer programs, $\forall\exists$ sentences,” in *DIMACS Series in Discrete Mathematics and Theoretical Computer Science Volume 1*, pp. 39–47, 1990.
- [72] KANNAN, R., “Algorithmic geometry of numbers,” *Annual Review of Comp. Sci.*, vol. 2, pp. 231–267, 1987.
- [73] KHACHIYAN, L. G., “A polynomial algorithm in linear programming,” *Soviet Mathematics Doklady*, vol. 20, pp. 191–194, 1979.
- [74] KHACHIYAN, L. G., “Polynomial algorithms in linear programming,” *USSR Computational Mathematics and Mathematical Physics*, vol. 20, pp. 53–72, 1980.
- [75] KHOT, S., “Hardness of approximating the shortest vector problem in lattices,” *J. ACM*, vol. 52, no. 5, pp. 789–808, 2005. Preliminary version in FOCS 2003.
- [76] KINCHINE, A., “A quantitative formulation of kronecker’s theory of approximation,” *Izv. Acad. Nauk SSSR*, vol. 12, pp. 113–122, 1948.
- [77] KLARTAG, B., “On convex perturbations with a bounded isotropic constant,” *Geom. and Funct. Anal.*, vol. 16(6), pp. 1274–1290, 2006.
- [78] KLARTAG, B., “On convex perturbations with a bounded isotropic constant,” *Geometric And Functional Analysis*, vol. 16, pp. 1274–1290, 2006.
- [79] KUPERBERG, G., “From the mahler conjecture to gauss linking integrals,” *Geometric And Functional Analysis*, vol. 18, pp. 870–892, 2008.
- [80] LAGARIAS, J. C., JR., H. W. L., and SCHNORR, C.-P., “Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice,” *Combinatorica*, vol. 10, no. 4, pp. 333–348, 1990.
- [81] LATALA, R. and OLESZKIEWICZ, K., “Gaussian measures of dilatations of convex symmetric sets,” *Annals of Probability*, vol. 27, no. 4, pp. 1922–1938, 1999.
- [82] LENSTRA, A. K., LENSTRA, H. W., and LOVÁSZ, L., “Factoring polynomials with rational coefficients,” *Mathematische Annalen*, vol. 261, pp. 515–534, 1982.
- [83] LENSTRA, A. K., LENSTRA, JR., H. W., and LOVÁSZ, L., “Factoring polynomials with rational coefficients,” *Mathematische Annalen*, vol. 261, pp. 515–534, December 1982.
- [84] LENSTRA, H. W., “Integer programming with a fixed number of variables,” *Mathematics of Operations Research*, vol. 8, pp. 538–548, November 1983.
- [85] LEWIS, D. R., “Ellipsoids defined by Banach ideal norms,” *Mathematika*, vol. 26, no. 1, pp. 18–29, 1979.

- [86] LOVÁSZ, L. and VEMPALA, S., “Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization,” in *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, (Washington, DC, USA), pp. 57–68, IEEE Computer Society, 2006.
- [87] LOVÁSZ, L. and VEMPALA, S., “Hit-and-run from a corner,” *SIAM J. Computing*, vol. 35, pp. 985–1005, 2006.
- [88] LOVÁSZ, L. and VEMPALA, S., “Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm,” *J. Comput. Syst. Sci.*, vol. 72, no. 2, pp. 392–417, 2006.
- [89] MICCIANCIO, D., “The shortest vector in a lattice is hard to approximate to within some constant,” *SIAM J. Comput.*, vol. 30, no. 6, pp. 2008–2035, 2000. Preliminary version in FOCS 1998.
- [90] MICCIANCIO, D. and REGEV, O., “Worst-case to average-case reductions based on Gaussian measures,” *SIAM J. Comput.*, vol. 37, no. 1, pp. 267–302, 2007. Preliminary version in FOCS 2004.
- [91] MICCIANCIO, D. and VOULGARIS, P., “A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations,” in *STOC*, pp. 351–358, 2010.
- [92] MILMAN, V., “Inegalites de brunn-minkowski inverse et applications at la theorie locales des espaces normes,” *C. R. Acad. Sci. Paris*, vol. 302, no. 1, pp. 25–28, 1986.
- [93] MILMAN, V., “Isomorphic symmetrization and geometric inequalities,” in *Geometric Aspects of Functional Analysis* (LINDENSTRAUSS, J. and MILMAN, V., eds.), vol. 1317 of *Lecture Notes in Mathematics*, pp. 107–131, Springer Berlin / Heidelberg, 1988.
- [94] MILMAN, V. and PAJOR, A., “Isotropic position and inertia ellipsoids and zonoids of the unit ball of a normed n -dimensional space,” *Geometric Aspects of Functional Analysis*, pp. 64–104, 1989.
- [95] MILMAN, V. and PAJOR, A., “Entropy and asymptotic geometry of non-symmetric convex bodies,” *Advances in Mathematics*, vol. 152, no. 2, pp. 314 – 335, 2000.
- [96] MILMAN, V. D., BOURGAIN, J., and KLARTAG, B., “Symmetrization and isotropic constants of convex bodies,” in *Geometric Aspects of Functional Analysis*, vol. 1850 of *Lecture Notes in Mathematics*, pp. 101–115, Springer Berlin / Heidelberg, 2004.
- [97] MINKOWSKI, H., *Geometrie Der Zahlen*. Leipzig and Berlin: R. G. Teubner, 1910.

- [98] NAZAROV, F., “The Hörmander proof of the Bourgain-Milman theorem.” Preprint, 2009.
- [99] NGUYEN, P. Q. and STERN, J., “The two faces of lattices in cryptology,” in *CaLC*, pp. 146–180, 2001.
- [100] NIVEN, I. M., *Diophantine approximations*. New York: Interscience Publishers, 1963.
- [101] ODLYZKO, A. M., “The rise and fall of knapsack cryptosystems,” in *Cryptography and Computational Number Theory* (POMERANCE, C., ed.), vol. 42 of *Proceedings of Symposia in Applied Mathematics*, pp. 75–88, 1990.
- [102] PEIKERT, C., “An efficient and parallel Gaussian sampler for lattices,” in *CRYPTO*, pp. 80–97, 2010.
- [103] PINSKY, M., *Introduction to Fourier Analysis and Wavelets*. American Mathematical Society, 2002.
- [104] PISIER, G., “Remarques sur un resultat non publie de b. maurey,” (*French*) *Sminaire d’Analyse Fonctionnelle 1980-81*, vol. Exp. No. V, 13, 1981.
- [105] PISIER, G., “Holomorphic semi-groups and the geometry of banach spaces,” *Annals of Mathematics*, vol. 115, pp. 375–392, 1982.
- [106] PISIER, G., *The Volume of Convex Bodies and Banach Space Geometry*. Cambridge University Press, 1989.
- [107] R. IMPAGLIAZZO, R. P., “The complexity of k -sat,” in *IEEE Conference on Computational Complexity*, 1999.
- [108] REGEV, O., “On lattices, learning with errors, random linear codes, and cryptography,” *J. ACM*, vol. 56, no. 6, pp. 1–40, 2009. Preliminary version in STOC 2005.
- [109] REGEV, O. and ROSEN, R., “Lattice problems and norm embeddings,” in *STOC*, pp. 447–456, 2006.
- [110] ROCKAFELLAR, R., *Convex Analysis*. Princeton, New Jersey: Princeton University Press, 2nd ed. ed., 1970.
- [111] ROGERS, C. and SHEPHARD, G., “The difference body of a convex body,” *Arch. Math.*, vol. 8, pp. 220–233, 1957.
- [112] ROGERS, C. and SHEPHARD, G., “Convex bodies associated with a given convex body,” *J. London Soc.*, vol. 33, pp. 270–281, 1958.
- [113] RUDELSON, M., “Distance between non-symmetric convex bodies and the MM^* -estimate,” *Positivity*, vol. 4, no. 8, pp. 161–178, 2000.

- [114] RUDIN, W., *Real and Complex Analysis*. New York, NY: McGraw-Hill, 2nd ed., 1974.
- [115] RUDIN, W., *Principles of Mathematical Analysis*. New York, NY: McGraw-Hill, 3rd ed., 1976.
- [116] SANTALÓ, L. A., “Un invariante afin para los cuerpos convexos del espacio de n dimensiones,” *Portugaliae Math.*, vol. 8, pp. 155–161, 1949.
- [117] SCHNEIDER, R., *Convex Bodies: The Brunn-Minkowski Theory*. New York: Cambridge University Press, 1993.
- [118] SCHNORR, C.-P., “A hierarchy of polynomial time lattice basis reduction algorithms,” *Theor. Comput. Sci.*, vol. 53, pp. 201–224, 1987.
- [119] SCHRIJVER, A., “On cutting planes,” *Annals of Discrete Mathematics*, vol. 9, pp. 291–296, 1980. *Combinatorics 79* (Proc. Colloq., Univ. Montréal, Montreal, Que., 1979), Part II.
- [120] SCHRIJVER, A., *Theory of Linear and Integer Programming*. New York, NY: Wiley-Interscience, 1986.
- [121] SHARIPOV, R., *Course of Linear Algebra and Multidimensional Geometry*. Publ. of Bashkir State University, 1996.
- [122] SHOR, N., “Cut-off method with space extension in convex programming problems,” *Kibernetika*, vol. 13, no. 1, pp. 94–95, 1977.
- [123] SIPSER, M., *Introduction to the Theory of Computation*. PWS Pub. Co., 1996.
- [124] SONNEVEND, G., “Applications of analytic centers for the numerical solution of semi-infinite convex programs arising in control theory,” in *System Modelling and Optimization* (SEBASTIAN, H. and TAMMER, K., eds.), vol. 143 of *Lecture Notes in Control and Information Sciences*, pp. 413–422, Springer Berlin / Heidelberg, 1990.
- [125] VAN EMDE BOAS, P., “Another NP-complete problem and the complexity of computing short vectors in a lattice,” Tech. Rep. 81-04, University of Amsterdam, 1981.
- [126] VEMPALA, S., “Recent progress and open problems in algorithmic convex geometry,” in *FSTTCS*, pp. 42–64, 2010.
- [127] YUDIN, D. B. and NEMIROVSKI, A. S., “Evaluation of the information complexity of mathematical programming problems (in russian),” *Ekonomika i Matematicheskie Metody*, vol. 13, no. 2, pp. 3–45, 1976.

VITA

Daniel Dadush is a French and American national born in London, England in 1982. His family moved to the United States in 1992 to the quiet suburbs of Washington D.C. in Potomac, Maryland. There he attended a French international school and transferred to an American public school for high school. Before matriculating at Brown University for college, he spent a year and a half in Milan, Italy working in a startup internet marketing firm. In 2006, he earned a S.C.B. in mathematics from Brown, where he graduated magna cum laude, received the David Howell premium for excellence in mathematics, and earned membership in Phi Beta Kappa. After completing his undergraduate degree, he spent six months interning at ITA software in Boston, and after applying for graduate school, traveled through India, Kazakhstan and China for five months before beginning his graduate studies. He began his PhD in Algorithms, Combinatorics and Optimization at Georgia Tech in 2007, where he received the Georgia Tech Institute Fellowship (GTIF), the most prestigious fellowship offered by the university. His research has been supported by the Achievement Rewards for College Scientists (ARCS) foundation and the Algorithms and Randomness Center (ARC) student fellowship. His work on cutting planes was awarded the INFORMS Optimization Society student paper prize in 2011. In recognition of his leadership within the ACO program, he was awarded the ACO outstanding student prize in 2012.