

Generalized compact knapsacks, cyclic lattices, and efficient one-way functions*

Daniele Micciancio
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0404, USA
daniele@cs.ucsd.edu

Abstract

We investigate the average-case complexity of a generalization of the compact knapsack problem to arbitrary rings: given m (random) ring elements $a_1, \dots, a_m \in R$ and a (random) target value $b \in R$, find coefficients $x_1, \dots, x_m \in S$ (where S is an appropriately chosen subset of R) such that $\sum a_i \cdot x_i = b$. We consider compact versions of the generalized knapsack where the set S is large and the number of weights m is small. Most variants of this problem considered in the past (e.g., when $R = \mathbb{Z}$ is the ring of the integers) can be easily solved in polynomial time even in the worst case. We propose a new choice of the ring R and subset S that yields generalized compact knapsacks that are seemingly very hard to solve on the average, even for very small values of m . Namely, we prove that for any unbounded function $m = \omega(1)$ with arbitrarily slow growth rate, solving our generalized compact knapsack problems *on the average* is at least as hard as the *worst-case* instance of various approximation problems over cyclic lattices. Specific worst-case lattice problems considered in this paper are the shortest independent vector problem SIVP and the guaranteed distance decoding problem GDD (a variant of the closest vector problem, CVP) for approximation factors $n^{1+\epsilon}$ almost linear in the dimension of the lattice.

Our results yield very efficient and provably secure one-way functions (based on worst-case complexity assumptions) with key size and time complexity almost linear in the security parameter n . Previous constructions with similar security guarantees required quadratic key size and computation time. Our results can also be formulated as a connection between the worst-case and average-case complexity of various lattice problems over cyclic and quasi-cyclic lattices.

Keywords: Knapsack problem, cyclic lattices, average-case complexity, one-way functions.

1 Introduction

Few problems in the theory of computational complexity and its application to the foundations of cryptography have been as controversial as the knapsack problem and its many variants, including the notorious NP-hard subset-sum problem [31]. The initial enthusiasm generated by the subset-sum based cryptosystem of Merkle and Hellman [40] in the late 70's was immediately followed by intensive cryptanalytic efforts that culminated in the early 80's with the total break of the system in its basic [65] and iterated version [9]. Still, the possibility of building cryptographic functions based on NP-hard problems, and the relatively high speed at which numbers can be added up (compared to modular multiplication and exponentiation

*An edited version of this paper appears in *Computational Complexity* 16(4):365-411 (2007). This is the author's copy of the paper. A preliminary version of this work appeared in the Proceedings of the 43rd Annual Symposium on Foundations of Computer Science - FOCS 2002, IEEE, with the title "Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions". Research supported in part by NSF grants CCR-0093029 and CCF-0634909, and a Sloan Research Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

operations required by number theoretic functions), prompted many researchers to suggest variants, fixes, and improvements (e.g., [22, 11]) to the initial Merkle-Hellman proposal. These efforts, which lasted for more than a decade, were invariably followed by attacks (e.g., [29, 64, 51, 55]) that seriously questioned the security of the systems either in theory or in practice. Recently, knapsack-like cryptographic functions have started attracting again considerable attention after Ajtai’s discovery [2] that the generalized subset-sum problem (over the additive group \mathbb{Z}_p^n of n -dimensional vectors modulo p) is provably hard to solve on the average based on a worst-case intractability assumption about certain lattice approximation problems for which no polynomial time solution is known. Following [2], Ajtai and Dwork [3] also proposed a public-key cryptosystem with similar security properties. But, unfortunately, even this proposal with strong theoretical security guarantees has been subject to practical attacks [52].

Attacks to subset-sum (or more generally knapsack) problems can be classified into two broad categories:

1. attacks targeted to specific public-key cryptosystems that try to exploit the special structure resulting from the embedding of a decryption trapdoor (e.g., [65]); and
2. attacks to generic subset-sum or knapsack instances that can be applied regardless of the existence of a trapdoor (e.g., [33, 12]).

The first class of attacks is usually stronger, meaning that it gives asymptotically good algorithms that succeed (with high probability) regardless of the value of the security parameter, but only applies to specific public-key cryptosystems whose underlying knapsack problems are not as hard as the general case. The second class of attacks is more general but only heuristic: the asymptotic complexity of these attacks is usually exponential, or their success rate negligible as a function of the security parameter. These methods are evaluated experimentally by testing them on specific problem instances (e.g., challenges or randomly generated ciphertexts) for typical values of the security parameter, and attacks can be usually avoided setting the security parameter to a sufficiently large value. Still, the effectiveness of these attacks, even for moderately large values of the security parameter, is currently considered the main practical obstacle to the design of cryptographic functions based on variants of the knapsack problem.

It is important to realize that the second class of attacks dismisses most knapsack cryptographic functions as practical alternatives to number theory based functions, not on the grounds of their inherent insecurity, but simply because of the large key sizes required to avoid heuristic attacks. In fact (especially if one drops the more ambitious goal of designing a public-key cryptosystem, and more modestly attempts to design cryptographic primitives with no trapdoors, like pseudo-random generators or one-way hash functions, etc.) there is theoretical evidence [27, 2, 3, 59] that subset-sum can indeed be a good source of computational hardness, at least from an asymptotic point of view. The main issue affecting the practical security of knapsack functions is efficiency. In a typical knapsack function, the key (corresponding to security parameter n) consists of $\Omega(n)$ numbers, each of which is n bits long. Therefore, the size of the resulting cryptographic key grows as $\Omega(n^2)$. Even if all known attacks to knapsack have exponential time complexity, one needs to set n to at least a few hundreds to make heuristic approaches (most notably lattice basis reduction [35, 63, 30, 53, 54]) ineffective or too costly. As a consequence, the resulting key can easily reach megabit sizes still without achieving a sufficient degree of security. Even if knapsack functions can still be competitive from a running time point of view, these huge key sizes are considered too big for most practical applications.

Generalized compact knapsacks. The impact of space efficiency on the practical security of knapsack based functions has long been recognized, even before the development of ingenious lattice-based attacks. A simple improvement that comes to mind is to use a so called *compact* knapsack: instead of using 0–1 combinations of $\Omega(n)$ input weights (resulting in $\Omega(n^2)$ key size), consider a smaller (constant, or slowly increasing) number of weights a_1, \dots, a_m and combine them with coefficients from a larger set, e.g., $\{0, \dots, 2^{\delta n}\}$ for some small constant $\delta > 0$. Notice that if $\delta = 0$, then we get the usual subset-sum problem, which can be solved (for $m = O(\log n)$) in polynomial time using exhaustive search. However, if $\delta = \Omega(1)$ then the search space becomes exponentially large, and exhaustive search is infeasible. Suggestions of this type appear already in Merkle and Hellman’s original paper [40] and subsequent works as a method to increase the bandwidth of the scheme. These early attempts to reduce the key size of knapsack based functions were subject to attacks

even more devastating than the general case: in [5] it is observed that the problem easily reduces to an integer programming instance with $O(m)$ variables, and therefore it can be solved in polynomial time for any constant value of $m(n) = O(1)$, or even any slowly growing function $m(n) = O(\log n / \log \log n)$. Attempts to use compact knapsacks to design efficient cryptographic functions persisted during the 90's [56, 36], but were always followed by cryptanalytic attacks [60, 13, 34].

In this paper we introduce and study a new class of compact knapsacks which are both very efficient and provably hard to solve in a strong sense similar to Ajtai's function [2]. The one-way function proposed in [2] can be described as a generalization of the integer knapsack problem to arbitrary rings. Specifically, for any ring R and subset $S \subset R$, consider the following problem: given ring elements $a_1, \dots, a_m \in R$ and a target value $b \in R$, find coefficients $x_i \in S$ such that $\sum_{i=1}^m a_i \cdot x_i = b$, where all operations are performed in the ring. In Ajtai's work, R is the product ring¹ \mathbb{Z}_p^n of n -dimensional vectors modulo p (for some polynomially bounded $p(n) = n^{O(1)}$) and $S = \{\mathbf{0}, \mathbf{1}\}$ consists of the additive and multiplicative identities of the ring. In particular, S has size 2, and the problem can be solved by exhaustive search in polynomial time when $m = O(\log n)$.

In this paper we study compact versions of the generalized knapsack problem, where the set S has size much larger than 2, so that exhaustive search is infeasible even for very small values of m . In the case of the ring \mathbb{Z}_p^n , the first idea that comes to mind is to use, as coefficients the set $S = \{0, 1\}^n$ of all binary vectors, or, more generally, the set $S = \{0, \dots, \lfloor p^\delta \rfloor\}^n$ of n -dimensional vectors with entries much smaller than p . Unfortunately, as for the case of the integer compact knapsack problem described above, this straightforward construction admits much faster solutions than exhaustive search: the resulting generalized compact knapsack is equivalent to n independent instances of the knapsack problem modulo p , which can be efficiently solved in the worst case for any polynomially bounded $p(n) = n^{O(1)}$ by dynamic programming, and on the average for $p(n) = n^{O(\log n)}$ using the methods in [17, 37].

Our contribution. The main contribution of this paper is the study of a new class of compact knapsack functions $f_{\mathbf{a}}(\mathbf{x}) = \sum_i a_i \cdot x_i$ that are provably hard to invert in a very strong sense, even when the number m of weights is very small. In particular, we prove that, for appropriate choice of ring R and subset $S \subset R$, and for any unbounded function $m(n) = \omega(1)$ (with arbitrarily slow growth rate) the compact knapsack function is at least as hard to invert *on the average* (even with non-negligible probability) as the *worst-case* instance of various lattice problems (for the special class of cyclic lattices, i.e., lattices that are invariant under cyclic rotations of the coordinates) for which no polynomial time algorithm is known.

Our generalized knapsack problems are defined by the ring $R = \mathbb{Z}_p^n$ of n -dimensional vectors modulo a prime p with the componentwise addition and *convolution product* operations. As in the previously discussed compact variant of Ajtai's function, the set $S = \{0, \dots, \lfloor p^\delta \rfloor\}^n$ consists of all n -dimensional vectors with small entries. Remarkably, using the convolution product operation (as opposed to componentwise multiplication) makes the problem considerably harder: solving random instances of our generalized compact knapsacks with non-negligible probability is as hard as approximating the shortest independent vector problem (as well as various other lattice problems) on cyclic lattices in the *worst case* within factors $n^{1+\epsilon}$ (for any $\epsilon > 0$) almost linear in the dimension of the lattice.

This results in strong one-way functions with average-case security guarantees based on a worst-case intractability assumption similar to Ajtai's function [2] (and subsequent improvements [10, 46, 48],) but with a much smaller key size $O(m \log p^n) = \omega(1) \cdot n \log n$, where $\omega(1)$ is an unbounded function with arbitrarily slow growth rate. (For comparison, [2, 10, 46, 48] require $m(n) = \Omega(n \log n)$, and key size $\Omega(n^2 \log^2 n)$.)

Our compact knapsack functions are also extremely fast, as, for appropriate choice of the parameters, they can be computed in almost linear time $O(n \log^c n)$ using the fast Fourier transform (FFT) in the evaluation of the convolution products. Specifically, the cost of evaluating our functions is equivalent to computing an almost constant number $\omega(1)$ of FFT operations on n -dimensional vectors modulo a small prime $p = n^{O(1)}$. The almost linear time evaluation algorithm together with the substantially smaller key size, make our generalized compact knapsack function even much faster than the already attractive subset-sum function.

¹The product ring R^n is the set of n -tuples with entries in R , with the componentwise addition and multiplication operations.

In the process of establishing our hardness result, we prove various properties of our knapsack functions that might be of independent interest. In particular, we prove that our compact knapsack function $f_{\mathbf{a}}(\mathbf{x})$ has very small collision probability, when the input is chosen uniformly at random. By a result of Rackoff (reported in [28]), this is enough to guarantee that the value $f_{\mathbf{a}}(\mathbf{x})$ (for randomly chosen \mathbf{a} and \mathbf{x}) is almost uniformly distributed over \mathbb{Z}_p^n and independent from $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_m)$. Moreover, this is true for arbitrary small values of $m(n) = \omega(1)$. Previous results of this kind for the subset-sum function relied on the additive structure of \mathbb{Z}_p^n alone, and required $m = \Omega(n \log p)$. Our proof makes substantial use of the multiplicative structure of the ring \mathbb{Z}_p^n (with the convolution product operation) and the characterization of its ideals as polynomial quotient rings.

Beside the technical contribution of a very efficient and provably secure one-way function based on a worst-case complexity assumption, we view the following as additional contributions of this paper: the introduction of the class of cyclic lattices as a source of interesting computational problems; casting a new light on the complexity of the compact knapsack problem showing that if the ring is appropriately chosen the problem can be substantially harder than the integer case; and demonstrating that the techniques initially developed in [2, 48] can be useful to study seemingly different problems, and still produce the same kind of strong worst-case/average-case security guarantees. In our view all these contributions are important steps toward the development of cryptographic functions that are both efficient and provably secure in a very strong sense. Finally, we remark that the problem of inverting our generalized compact knapsack function can be equivalently formulated as the problem of finding a lattice point (in a quasi-cyclic² lattice) close to a given target. (See Section 5 for details.) Therefore, our main result is interesting also from a purely complexity theoretic perspective, since it establishes a connection between the worst-case and average-case complexity of solving various lattice problems on (quasi-)cyclic lattices. This is analogous to previous results [2, 10, 19, 46, 48] connecting the worst-case and average-case complexity of problems on arbitrary lattices, but adapted to the special class of lattices with (quasi-)cyclic structure.

Related work. The first construction of one-way function that is provably secure based on a worst-case complexity assumption was given by Ajtai in [2]. Subsequent work [10, 46, 48] focused on weakening the required worst-case complexity assumption. In this paper, the goal is to improve the efficiency of the one-way function.

This paper is an almost complete rewriting and substantial improvement of an extended abstract [44] presented at FOCS 2002. In particular, in [44] the author proved that solving the generalized compact knapsack on the average when $m = O(\log n)$ is at least as hard as approximating various lattice problems in the worst case within a factor $n^{3+\epsilon}$. Here, we prove a new regularity theorem³ for compact knapsack functions (Theorem 4.2) and incorporate the recently developed Gaussian distribution techniques of [48] to obtain an improved result that holds for any function $m = \omega(1)$ with arbitrarily slow growth rate, and worst-case approximation factors $n^{1+\epsilon}$, almost linear in the dimension of the lattice.

Following the writing of this paper, it has been shown [38, 57] that variants of our generalized compact knapsack function are not only one way, but also collision resistant, a stronger and very useful cryptographic property. These improvements and related open problems are discussed in Section 5.

From a theoretical point of view, the main difference between our one-way functions and those studied in previous work (e.g., [2, 48] and related papers,) is that our functions are based on the worst-case intractability of lattice problems on a class of lattices with a special cyclic structure. Many lattice problems are known to be NP-hard even in their approximation versions for sufficiently small approximation factors. For example, the shortest vector problem (SVP) is NP-hard (under randomized reductions) to approximate within any constant factor [1, 43, 32], while the closest vector problem (CVP) is NP-hard to approximate even within quasi polynomial factors $n^{O(1/\log \log n)}$ [67, 6, 15]. These results support the conjecture that lattice problems are hard to solve in the worst case, at least for arbitrary lattices. It is natural to ask whether lattice problems remain hard even when the input lattice is cyclic.

Very little is known about the computational complexity of lattice problems on cyclic lattices. In fact, as

²A lattice is called quasi-cyclic if it is invariant under rotations of the coordinates by a number of positions possibly greater than 1.

far as we know, cyclic lattices have received little or no attention so far. From an algorithmic point of view, it is not clear how to exploit the cyclic structure of the lattice in state of the art lattice algorithms, e.g., lattice basis reduction. The only algorithmic results related to cyclic lattices we are aware of are [39, 26, 66, 18]. The first paper [39] shows how the solution of certain lattice problems can be speeded up by a factor n when the lattice is cyclic of dimension n . This is a quite modest improvement since the running time of the best algorithms to solve these problems over general lattices is exponential in n . A more interesting algorithmic result is given in [26, 66, 18]. The problem considered in [26] (and solved building on previous algorithms from [66, 18]) is the following: given the autocorrelation³ of a vector \mathbf{x} , retrieve \mathbf{x} . This problem (which arises from applications in n -dimensional crystallography) is related to cyclic lattices by the fact that the autocorrelation of \mathbf{x} can be expressed as a vector in the cyclic lattice generated by \mathbf{x} . This problem is quite different from the worst-case computational problems on cyclic lattices considered in this paper, and it is not clear if the techniques of [26, 66, 18] can be used to speed up the solution of other problems, like SVP, CVP or their variants SIVP (shortest independent vector problem) and GDD (generalized distance decoding) over cyclic lattices. Based on the current state of knowledge, it seems reasonable to conjecture that approximation problems on cyclic lattices are computationally hard, at least in the worst case and for small polynomial approximation factors. In order to further support this conjecture, it would be nice to prove NP-hardness results for lattice problems when restricted to cyclic lattices.

We remark that our definition of cyclic lattices is analogous to the definition of cyclic codes, one of the most useful and widely studied classes of codes in coding theory. Still, no polynomial time algorithm is known for many computational problems on cyclic codes (or lattices). A very recent result somehow suggesting that no such polynomial time algorithm may exist is the proof in [23] that the nearest codeword problem (the coding analogue of the closest vector problem for lattices) for appropriately shortened Reed-Solomon codes is NP-hard. Reed-Solomon codes are a well known class of cyclic codes, so the result in [23] seems to suggest that the nearest codeword problem is hard even when the code is cyclic. Unfortunately, shortening the Reed-Solomon code (as done in [23]) destroys the cyclic structure of the code, so, the results in [23] do not imply the NP-hardness of the nearest codeword problem over cyclic codes. We leave, as an open problem, to prove hardness results for any lattice or coding problem over cyclic lattices or codes. Is the shortest vector problem on cyclic lattices NP-hard? Is the shortest independent vector problem on cyclic lattices NP-hard? What about the closest vector problem on cyclic lattices? Is the closest vector problem NP-hard even for fixed families of cyclic lattices as shown (for arbitrary lattices) in [41, 16, 58]?

Organization The rest of the paper is organized as follows. In Section 2 we recall basic notation, definitions and results needed in this paper. In Section 3 we prove two preliminary lemmas about cyclic lattices that will be used in the proof of our main result. In Section 4 we present the main technical result of the paper: we formally define our generalized compact knapsack function, and prove that inverting the function on the average is at least as hard as the worst-case instance of various lattice problems on cyclic lattices. In the process, we also establish various other properties of our compact knapsack function that might be of independent interest, e.g., we bound the collision probability of the function, and prove that the function is almost regular. Section 5 concludes with a discussion additional related results and open problems.

2 Preliminaries

In this section we introduce some notational conventions, and recall basic definitions and results about the statistical distance, hash functions, lattices and Gaussian probability distributions.

For any real $r \geq 0$, $[r]$ denotes the set $\{0, \dots, [r]\}$ of all non-negative integers not greater than r . The uniform probability distribution over a set S is denoted $U(S)$. We use the standard asymptotic notation $f = O(g)$ (or $g = \Omega(f)$) when $\limsup_{n \rightarrow \infty} |f(n)/g(n)| < \infty$, $f = o(g)$ (or $g = \omega(f)$) when $\lim_{n \rightarrow \infty} |f(n)/g(n)| = 0$, and $f = \Theta(g)$ when $f = O(g)$ and $f = \Omega(g)$. A function $f(n)$ is negligible (denoted $f(n) = n^{-\omega(1)}$) if for every c there exists an n_0 such that $|f(n)| < 1/n^c$ for all $n > n_0$.

³The autocorrelation of a vector \mathbf{x} is the convolution of \mathbf{x} with itself $\mathbf{x} \otimes \mathbf{x}$. See Section 2 for a definition of the convolution product \otimes .

2.1 Statistical distance

The statistical distance is a measure of how two probability distributions are far apart from each other, and it is a convenient tool in the analysis of randomized algorithms and reductions. In this subsection we define the statistical distance and state some simple facts that will be used in the analysis of the reductions in this paper. All the properties of the statistical distance stated in this subsection are easily verified. For more details the reader is referred to [47, Chapter 8].

Definition 2.1 *Let X and Y be two discrete random variables over a (countable) set A . The statistical distance between X and Y is the quantity*

$$\Delta(X, Y) = \frac{1}{2} \sum_{a \in A} |\Pr\{X = a\} - \Pr\{Y = a\}|.$$

In the case of continuous random variables, the statistical distance between X and Y is

$$\Delta(X, Y) = \frac{1}{2} \int_A |\delta_X(a) - \delta_Y(a)| da,$$

where δ_X and δ_Y are the probability density functions of X and Y respectively.

We say that two random variables X, Y are identically distributed (written $X \equiv Y$) if and only if $\Pr\{X \in S\} = \Pr\{Y \in S\}$ for every $S \subseteq A$. The reader can easily check that the statistical distance satisfies the usual properties of distance functions, i.e., $\Delta(X, Y) \geq 0$ (with equality if and only if $X \equiv Y$), $\Delta(X, Y) = \Delta(Y, X)$, and $\Delta(X, Z) \leq \Delta(X, Y) + \Delta(Y, Z)$.

The following proposition shows that applying a (possibly randomized) function to two distributions does not increase the statistical distance.

Proposition 2.2 *Let X, Y be two random variables taking values in a common set A . For any (possibly randomized) function f with domain A , the statistical distance between $f(X)$ and $f(Y)$ is at most*

$$\Delta(f(X), f(Y)) \leq \Delta(X, Y). \tag{2.1}$$

As a corollary, we easily obtain the following.

Corollary 2.3 *If X and Y are random variables over set A and $p: A \rightarrow \{0, 1\}$ is a predicate, then*

$$|\Pr\{p(X) = 1\} - \Pr\{p(Y) = 1\}| \leq \Delta(X, Y). \tag{2.2}$$

Another useful property of the statistical distance is the following.

Proposition 2.4 *Let X_1, \dots, X_k and Y_1, \dots, Y_k be two lists of totally independent random variables. Then*

$$\Delta((X_1, \dots, X_k), (Y_1, \dots, Y_k)) \leq \sum_{i=1}^k \Delta(X_i, Y_i). \tag{2.3}$$

2.2 One-way hash function families

A function family $\{f_a: X \rightarrow R\}_{a \in A}$ is a collection of functions (indexed by a set of keys A) with a common domain X and range R . A (polynomial) function ensemble is a sequence $\{f_a: X_n \rightarrow R_n\}_{a \in A_n}$ of function families (indexed by a security parameter $n \in \mathbb{N}$) such that $\log |A_n|$, $\log |X_n|$ and $\log |R_n|$ are all polynomial in n . We assume that the elements of the sets A_n , X_n and R_n can be efficiently represented with $\log_2 |A_n|$, $\log_2 |X_n|$ and $\log_2 |R_n|$ bits respectively, membership in the sets can be decided in polynomial time, and there is a probabilistic polynomial time algorithm to sample from those sets with (almost) uniform distribution. It is also common to assume that the functions f_a are efficiently computable, in the sense that there is a

polynomial time algorithm that on input $n, a \in A_n$ and $x \in X_n$, outputs $f_a(x)$. All function ensembles considered in this paper have these properties, namely the sets A_n, X_n, R_n have efficient representations and the functions f_a are efficiently computable.

A function (ensemble) is one-way if it is (easy to compute, but) computationally hard to invert, i.e., no algorithm can efficiently solve the following function inversion problem: given a pair $(a, r) \in A_n \times R_n$, find an $x \in X_n$ such that $f_a(x) = r$. One-wayness is an average-case complexity property, i.e., it requires that the function inversion problem is computationally hard when the input $(a, r) \in A_n \times R_n$ is selected at random. The exact definition, for the case of function ensembles, is given below.

Definition 2.5 *A function ensemble $\{f_a : X_n \rightarrow R_n\}_{a \in A_n}$ is one-way if for any probabilistic polynomial time algorithm \mathcal{A} , the probability that $f_a(\mathcal{A}(n, a, f_a(x))) = f_a(x)$ (when $a \in A_n$ and $x \in X_n$ are selected uniformly at random) is negligible in n .*

Notice that the input distribution underlying the definition of one-way function is not the uniform distribution over $A_n \times R_n$, but rather it corresponds to choosing the target value $r \in R_n$ as the image of a uniformly random solution $x \in X$. For any function ensemble $\mathcal{H} = \{f_a : X \rightarrow R\}_{a \in A}$, we write $\text{OWF}(\mathcal{H})$ to denote the probability distribution $\{(a, f_a(x)) : a \in A, x \in X_n\}$ underlying the definition of one-way function, and $U(A \times R)$ to denote the uniform probability distribution over $A \times R$. We remark that Definition 2.5 corresponds to the notion of *strong* one-way function, i.e., it is required that the success probability of any probabilistic polynomial time algorithm in solving the function inversion problem (when the input is chosen according to distribution $\text{OWF}(\mathcal{H})$) is negligible.

The function families $\mathcal{H} = \{f_a : X \rightarrow R\}_{a \in A}$ considered in this paper have the property that the input size $\log |X|$ is strictly bigger than the output size $\log |R|$, i.e., the functions “compress” the size of the input by a factor $\log |X| / \log |R|$. Such functions have many important applications in computer science and cryptography, and are generically called *hash* functions. In order to be useful, hash functions must satisfy some additional properties. A typical requirement is that if $a \in A$ and $x \in X$ are chosen uniformly at random, the distribution of $f_a(x) \in R$ is almost uniform and independent from a . In other words, $\text{OWF}(\mathcal{H})$ is statistically close to the uniform distribution $U(A \times R)$.

Definition 2.6 *Let $\mathcal{H} = \{f_a : X \rightarrow R\}_{a \in A}$ be a hash function family. We say that \mathcal{H} is ϵ -regular if the statistical distance between $\text{OWF}(\mathcal{H})$ and the uniform distribution $U(A \times R)$ is at most ϵ . A hash function ensemble $\{\mathcal{H}_n\}$ is called almost regular if there exists a negligible function $\epsilon(n) = n^{-\omega(1)}$ such that \mathcal{H}_n is $\epsilon(n)$ -regular for every n .*

We remark that if a function is ϵ -regular for $\epsilon = 0$, then the function maps the uniform input distribution to the uniform output distribution. So, Definition 2.6 is a generalization of the standard notion of regular function.

2.3 Lattices

Throughout the paper, we use column notation for all vectors, and use $(\cdot)^T$ to denote the matrix transposition operation. For example, $\mathbf{x} = (x_1, \dots, x_n)^T$ is the n -dimensional column vector with entries x_1, \dots, x_n , and $[\mathbf{x}, \dots, \mathbf{x}]$ is the $n \times n$ matrix with all columns equal to \mathbf{x} .

An n -dimensional *lattice*⁴ is the set of all integer combinations $\{\sum_{i=1}^n x_i \mathbf{b}_i : x_i \in \mathbb{Z}\}$ of n linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_n$ in \mathbb{R}^n . The set of vectors $\mathbf{b}_1, \dots, \mathbf{b}_n$ is called a *basis* for the lattice, and can be compactly represented by the matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{n \times n}$ having the basis vectors as columns. The lattice generated by \mathbf{B} is denoted $\mathcal{L}(\mathbf{B})$. Notice that $\mathcal{L}(\mathbf{B}) = \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^n\}$, where $\mathbf{B}\mathbf{x}$ is the usual matrix-vector multiplication. For any basis \mathbf{B} , we define the fundamental parallelepiped $\mathcal{P}(\mathbf{B}) = \{\mathbf{B}\mathbf{x} : \forall i. 0 \leq x_i < 1\}$. The following lemma shows how to sample lattice points uniformly at random from the fundamental parallelepiped associated to a given sublattice.

⁴For simplicity, in this paper we restrict all definitions to full dimensional lattices.

Lemma 2.7 ([47, Proposition 8.2]) *There is a probabilistic polynomial time algorithm that on input a lattice basis \mathbf{B} and a full rank sublattice $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$, outputs a lattice point $\mathbf{x} \in \mathcal{L}(\mathbf{B}) \cap \mathcal{P}(\mathbf{S})$ chosen uniformly at random.*

The dual of a lattice $\mathcal{L}(\mathbf{B})$ (denoted $\mathcal{L}(\mathbf{B})^*$) is the lattice generated by the matrix $(\mathbf{B}^{-1})^T$, and consists of all vectors that have integer scalar product with all lattice vectors.

For any vector $\mathbf{x} = (x_1, \dots, x_n)^T$, define the cyclic rotation $rot(\mathbf{x}) = (x_n, x_1, \dots, x_{n-1})^T$, and the corresponding circulant matrix $Rot(\mathbf{x}) = [\mathbf{x}, rot(\mathbf{x}), rot^2(\mathbf{x}), \dots, rot^{n-1}(\mathbf{x})]$. (Notice that \mathbf{x} , and $rot^i(\mathbf{x})$ are all column vectors, and $Rot(\mathbf{x})$ is the matrix whose columns are the cyclic rotations of \mathbf{x} by construction. It is easy to see that also the rows of $Rot(\mathbf{x})$ are all rotations of the same vector but with the entries in reverse order. For example, the last row of $Rot(\mathbf{x})$ is (x_n, \dots, x_1) .) A lattice $\mathcal{L}(\mathbf{B})$ is cyclic if it is closed under the rotation operation, i.e., if $\mathbf{x} \in \mathcal{L}(\mathbf{B})$ implies $rot(\mathbf{x}) \in \mathcal{L}(\mathbf{B})$. It is easy to see that a lattice is cyclic if and only if $\mathcal{L}(\mathbf{B}) = rot(\mathcal{L}(\mathbf{B}))$.

The convolution product of two vectors \mathbf{x} and \mathbf{y} is the vector

$$\mathbf{x} \otimes \mathbf{y} = Rot(\mathbf{x}) \cdot \mathbf{y} = \mathbf{x} \cdot y_1 + rot(\mathbf{x}) \cdot y_2 + \dots + rot^{n-1}(\mathbf{x}) \cdot y_n$$

with entries defined by the equation

$$(\mathbf{x} \otimes \mathbf{y})_k = \sum_{i+j=k+1 \pmod n} x_i \cdot y_j,$$

e.g., $(\mathbf{x} \otimes \mathbf{y})_n = x_n y_1 + x_{n-1} y_2 + \dots + x_1 y_n$. It can be easily verified that the convolution product is associative and commutative, i.e., it satisfies the equational axioms $\mathbf{x} \otimes (\mathbf{y} \otimes \mathbf{z}) = (\mathbf{x} \otimes \mathbf{y}) \otimes \mathbf{z}$, and $\mathbf{x} \otimes \mathbf{y} = \mathbf{y} \otimes \mathbf{x}$. Moreover, it distributes over the vector addition operation: $(\mathbf{x} + \mathbf{y}) \otimes \mathbf{z} = \mathbf{x} \otimes \mathbf{z} + \mathbf{y} \otimes \mathbf{z}$. Therefore, $(\mathbb{R}^n, +, \otimes)$ is a commutative ring with identity $\mathbf{e}_1 = (1, 0, \dots, 0)^T$.

The Euclidean norm of a vector \mathbf{x} is the quantity $\|\mathbf{x}\| = \sqrt{\sum_i x_i^2}$. Other norms used in this paper are the ℓ_1 norm $\|\mathbf{x}\|_1 = \sum_i |x_i|$ and the max norm $\|\mathbf{x}\|_\infty = \max_i |x_i|$. These norms and the convolution product are related by the following inequalities, valid for any n -dimensional vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$:

$$\|\mathbf{x}\| \leq \|\mathbf{x}\|_1 \leq \sqrt{n} \|\mathbf{x}\| \tag{2.4}$$

$$\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\| \leq \sqrt{n} \|\mathbf{x}\|_\infty \tag{2.5}$$

$$\|\mathbf{x} \otimes \mathbf{y}\|_\infty \leq \|\mathbf{x}\| \cdot \|\mathbf{y}\| \tag{2.6}$$

$$\|\mathbf{x} \otimes \mathbf{y}\|_\infty \leq \|\mathbf{x}\|_1 \cdot \|\mathbf{y}\|_\infty. \tag{2.7}$$

For any matrix or set of vectors \mathbf{S} , we denote by $\|\mathbf{S}\| = \max_i \|\mathbf{s}_i\|$ the norm of the longest (column) vector in \mathbf{S} .

The *minimum distance* of a lattice $\mathcal{L}(\mathbf{B})$, denoted $\lambda_1(\mathcal{L}(\mathbf{B}))$, is the minimum distance between any two (distinct) lattice points and equals the length of the shortest nonzero lattice vector:

$$\lambda_1(\mathcal{L}(\mathbf{B})) = \min\{\text{dist}(\mathbf{x}, \mathbf{y}) : \mathbf{x} \neq \mathbf{y} \in \mathcal{L}(\mathbf{B})\} = \min\{\|\mathbf{x}\| : \mathbf{x} \in \mathcal{L}(\mathbf{B}) \setminus \{\mathbf{0}\}\}.$$

The notion of minimum distance can be generalized to define the i th successive minimum λ_i as the smallest radius r such that the closed sphere $\bar{\mathcal{B}}(r) = \{\mathbf{x} : \|\mathbf{x}\| \leq r\}$ contains i linearly independent lattice points:

$$\lambda_i(\mathcal{L}(\mathbf{B})) = \min\{r : \dim(\text{span}(\mathcal{L}(\mathbf{B}) \cap \bar{\mathcal{B}}(r))) \geq i\}$$

Another important constant associated to a lattice is the covering radius. The covering radius $\rho(\mathcal{L}(\mathbf{B}))$ of a lattice is the maximum distance $\text{dist}(\mathbf{x}, \mathcal{L}(\mathbf{B}))$ when \mathbf{x} ranges over the entire space \mathbb{R}^n :

$$\rho(\mathcal{L}(\mathbf{B})) = \max\{\text{dist}(\mathbf{x}, \mathcal{L}(\mathbf{B})) : \mathbf{x} \in \mathbb{R}^n\}.$$

A sublattice of $\mathcal{L}(\mathbf{B})$ is a (full rank) lattice $\mathcal{L}(\mathbf{S})$ such that $\mathcal{L}(\mathbf{S}) \subseteq \mathcal{L}(\mathbf{B})$.

In many algorithmic problems on point lattices the quality of a solution is measured with respect to some specific lattice parameter, e.g., the length λ_1 of the shortest nonzero vector, or the radius λ_n of the smallest sphere containing n linearly independent lattice vectors. For example, the $\gamma(n)$ -approximate shortest vector problem asks to find a nonzero vector in a lattice $\mathcal{L}(\mathbf{B})$ of length at most $\gamma(n) \cdot \lambda_1(\mathcal{L}(\mathbf{B}))$, where n is the rank of the lattice. For technical reasons, in this paper we consider generalized versions of various lattice problems where the quality of the solution is measured with respect to an arbitrary function of the lattice $\phi(\mathcal{L}(\mathbf{B}))$. The first of these problems is the following generalization of the shortest independent vector problem introduced in [46].

Definition 2.8 *The generalized shortest independent vectors problem⁵ SIVP_γ^ϕ , given an n -dimensional lattice \mathbf{B} , asks for a set of n linearly independent lattice vectors $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{S}\| \leq \gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$.*

The shortest independent vectors problem SIVP_γ (studied in [8] and used in [2, 10, 46, 48] as a source of computational hardness) corresponds to SIVP_γ^ϕ with $\phi = \lambda_n$. Another problem that will play a fundamental role in this paper is the following.

Definition 2.9 *The guaranteed distance decoding problem (GDD_γ^ϕ), given a lattice \mathbf{B} and a target point $\mathbf{t} \in \text{span}(\mathbf{B})$, asks for a lattice point $\mathbf{x} \in \mathcal{L}(\mathbf{B})$ such that $\text{dist}(\mathbf{t}, \mathbf{x}) \leq \gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$, where n is the rank of the lattice.*

This time it is natural to set $\phi = \rho$ to the covering radius of the lattice, because for any lattice basis \mathbf{B} and target $\mathbf{t} \in \mathbb{R}^n$, there is always a lattice point within distance $\rho(\mathcal{L}(\mathbf{B}))$ from \mathbf{t} . When $\phi = \rho$, we omit the superscript, and simply write GDD_γ . GDD_γ is an interesting variant of the closest vector problem (CVP), where the quality of the solution is measured with respect to the worst possible distance $\max_{\mathbf{t} \in \mathbb{R}^n} \text{dist}(\mathbf{t}, \mathcal{L}(\mathbf{B}))$ rather than the distance of the given target $\text{dist}(\mathbf{t}, \mathcal{L}(\mathbf{B}))$.

No polynomial time algorithm to solve SIVP_γ or GDD_γ within polynomial approximation factors $\gamma(n) = n^{O(1)}$ is known. A well known polynomial time algorithm for approximating SIVP is the basis reduction algorithm of Lentra, Lentra, and Lovász [35], which on input a lattice \mathbf{B} , computes a so-called *LLL-reduced* basis \mathbf{S} for the same lattice. The exact definition of LLL-reduced basis is not important here. All we need in this paper is that the LLL-reduced basis satisfies $\|\mathbf{S}\| \leq 2^n \lambda_n(\mathcal{L}(\mathbf{B}))$, i.e., it solves SIVP_γ for approximation factors $\gamma(n) = 2^n$. A well known method to find lattice points close to a given target is Babai's nearest plane algorithm [7]. This is a polynomial time algorithm that on input a lattice \mathbf{S} and a target point \mathbf{t} , finds a lattice vector $\mathbf{x} \in \mathcal{L}(\mathbf{S})$ within distance $\|\mathbf{x} - \mathbf{t}\| \leq (\sqrt{n}/2)\|\mathbf{S}\|$ from the target. Notice that the quality of the solution depends on $\|\mathbf{S}\|$. For example, when used in conjunction with the LLL basis reduction algorithm, the nearest plane algorithm returns a lattice point within distance $\sqrt{n}2^{n-1}\lambda_n(\mathcal{L}(\mathbf{B})) \leq \sqrt{n}2^n\rho(\mathcal{L}(\mathbf{B}))$ from the target, i.e., it solves GDD_γ for approximation factor $\gamma(n) = \sqrt{n}2^n$. Slightly better approximations (namely, for slightly subexponential factors $2^{O(n \log \log n / \log n)}$) can be computed in (probabilistic) polynomial time using more complex algorithms [62, 4], but they offer no advantages in the context of our paper.

2.4 Gaussian distributions

We use the Gaussian distribution techniques recently introduced in [48] to simplify and improve the results described in a preliminary version of this paper [44]. In this subsection we recall all the required definitions and results from [48]. For any vectors \mathbf{c}, \mathbf{x} and any $s > 0$, let

$$\rho_{s,\mathbf{c}}(\mathbf{x}) = e^{-\pi\|(\mathbf{x}-\mathbf{c})/s\|^2}$$

be a Gaussian function centered in \mathbf{c} scaled by a factor of s . The total measure associated to $\rho_{s,\mathbf{c}}$ is $\int_{\mathbf{x} \in \mathbb{R}^n} \rho_{s,\mathbf{c}}(\mathbf{x}) d\mathbf{x} = s^n$. So, $\int_{\mathbf{x} \in \mathbb{R}^n} (\rho_{s,\mathbf{c}}(\mathbf{x})/s^n) d\mathbf{x} = 1$ and $\rho_{s,\mathbf{c}}/s^n$ is a probability density function. As noted in [48], $\rho_{s,\mathbf{c}}/s^n$ can be expressed as the sum of n orthogonal 1-dimensional Gaussian distributions, and

⁵In previous papers, this problem was denoted GIVP. Here we use the standard notation for the *shortest independent vector problem* SIVP , annotated with the superscript ϕ .

each of them can be efficiently approximated with arbitrary precision using standard techniques. So, the distribution $\rho_{s,\mathbf{c}}/s^n$ can be efficiently approximated. For simplicity, in this paper we work with real numbers and assume we can sample from $\rho_{s,\mathbf{c}}/s^n$ exactly. In practice, when only finite precision is available, $\rho_{s,\mathbf{c}}/s^n$ can be approximated by picking a fine grid, and selecting points from the grid with probability approximately proportional to $\rho_{s,\mathbf{c}}/s^n$. All our arguments can be made rigorous by selecting a sufficiently fine grid.

Functions are extended to sets in the usual way; e.g., $\rho_{s,\mathbf{c}}(A) = \sum_{\mathbf{x} \in A} \rho_{s,\mathbf{c}}(\mathbf{x})$ for any countable set A . For any s, \mathbf{c} and lattice Λ , define the discrete probability distribution (over the lattice Λ)

$$D_{\Lambda,s,\mathbf{c}}(\mathbf{x}) = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\Lambda)},$$

where $\mathbf{x} \in \Lambda$. Intuitively, $D_{\Lambda,s,\mathbf{c}}$ is the conditional probability⁶ that a random variable with probability density function ($\rho_{s,\mathbf{c}}/s^n$) takes the value \mathbf{x} given that the value of the random variable belongs to the lattice Λ . For brevity, we sometimes omit s or \mathbf{c} from the notation $\rho_{s,\mathbf{c}}$ and $D_{\Lambda,s,\mathbf{c}}$. When \mathbf{c} or s are not specified, we assume that they are the origin and 1 respectively.

In [48] Gaussian distributions are used to define a new lattice invariant, called the *smoothing parameter*, defined as follows.

Definition 2.10 *For an n -dimensional lattice Λ , and positive real $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest s such that $\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$.*

In [48] many important properties of the smoothing parameter are established. Here we only need the following three bounds. The first one shows that the smoothing parameter is the amount of Gaussian noise that needs to be added to a lattice in order to get an almost uniform distribution.

Lemma 2.11 ([48, Lemma 4.1]) *Let $\rho_s/s^n \bmod \mathbf{B}$ be the distribution obtained by sampling a point according to the probability density function ρ_s/s^n and reducing the result modulo \mathbf{B} . For any lattice $\mathcal{L}(\mathbf{B})$, the statistical distance between $\rho_s/s^n \bmod \mathbf{B}$ and the uniform distribution over $\mathcal{P}(\mathbf{B})$ is at most $\frac{1}{2}\rho_{1/s}(\mathcal{L}(\mathbf{B})^* \setminus \{\mathbf{0}\})$. In particular, if $s \geq \eta_\epsilon(\mathcal{L}(\mathbf{B}))$, then the distance $\Delta(\rho_s/s^n \bmod \mathbf{B}, U(\mathcal{P}(\mathbf{B})))$ is at most $\epsilon/2$.*

The second property shows that if s is sufficiently large, then the central second moment of the distribution $D_{\Lambda,s,\mathbf{c}}$ is essentially the same as the one of the continuous Gaussian distribution $\rho_{c,s}/s^n$.

Lemma 2.12 ([48, Lemma 4.2]) *For any n -dimensional lattice Λ , point $\mathbf{c} \in \mathbb{R}^n$, unit vector \mathbf{u} , and reals $0 < \epsilon < 1$, $s \geq 2\eta_\epsilon(\Lambda)$*

$$\left| E_{\mathbf{x} \sim D_{\Lambda,s,\mathbf{c}}} [\langle \mathbf{x} - \mathbf{c}, \mathbf{u} \rangle^2] - \frac{s^2}{2\pi} \right| \leq \frac{\epsilon s^2}{1 - \epsilon}.$$

The last property bounds the smoothing parameter in terms of λ_n .

Lemma 2.13 ([48, Lemma 3.3]) *For any n -dimensional lattice Λ and positive real $\epsilon > 0$,*

$$\eta_\epsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1 + 1/\epsilon))}{\pi}} \cdot \lambda_n(\Lambda).$$

In particular, for any super-logarithmic function $\omega(\log n)$ there is a negligible function $\epsilon(n)$ such that $\eta_\epsilon(\Lambda) \leq \sqrt{\omega(\log n)} \cdot \lambda_n$.

⁶We are conditioning on an event that has probability 0; this can be made rigorous by standard techniques.

3 Two lemmas about cyclic lattices

In this section we prove two preliminary lemmas about cyclic lattices that will be used in the proof of our main results in the next section. The results are presented here because their formulation is largely independent from the specific reduction in which they are used, and might be of independent interest.

The first lemma gives an efficient algorithm to select a (full rank) cyclic lattice generated by a single short vector \mathbf{c} from an arbitrary input lattice \mathbf{S} . We remark that the lemma below will be used in settings where the vectors in \mathbf{S} belong to a cyclic lattice $\mathcal{L}(\mathbf{B})$, so that the cyclic lattice $\mathcal{L}(\text{Rot}(\mathbf{c}))$ generated by $\mathbf{c} \in \mathcal{L}(\mathbf{S})$ is a sublattice of $\mathcal{L}(\mathbf{B})$. However, it is not generally the case that $\mathcal{L}(\text{Rot}(\mathbf{c}))$ is a sublattice of $\mathcal{L}(\mathbf{S})$ because $\mathcal{L}(\mathbf{S})$ may not be cyclic.

Lemma 3.1 *There exists a polynomial time algorithm that on input a full rank n -dimensional lattice \mathbf{S} , computes a vector $\mathbf{c} \in \mathcal{L}(\mathbf{S})$ such that $\|\mathbf{c}\|_1 \leq 2n \cdot \|\mathbf{S}\|$ and $\text{Rot}(\mathbf{c})$ has full rank.*

Proof: Let $S = \|\mathbf{S}\|$. We use Babai's nearest plane algorithm [7] to find a vector $\mathbf{c} \in \mathcal{L}(\mathbf{S})$ within Euclidean distance $(\sqrt{n}/2) \cdot S$ from $nS\mathbf{e}_1$. Notice that the ℓ_1 norm of \mathbf{c} is at most

$$\begin{aligned} \|\mathbf{c}\|_1 &\leq \|(nS \cdot \mathbf{e}_1)\|_1 + \|\mathbf{c} - nS\mathbf{e}_1\|_1 \\ &\leq nS + \sqrt{n}\|\mathbf{c} - nS\mathbf{e}_1\| \\ &\leq 1.5 \cdot nS. \end{aligned}$$

It remains to show that $\text{Rot}(\mathbf{c})$ is non-singular, or equivalently, the n -dimensional volume of $\mathcal{P}(\text{Rot}(\mathbf{c}))$ is nonzero. Notice that $\mathcal{P}(\text{Rot}(\mathbf{c}))$ is an almost cubic parallelepiped obtained by perturbing the main vertexes of a hypercube of size $l = nS$ by at most $\epsilon = (\sqrt{n}/2)S$. In [45] it is shown that, for all $\epsilon < \sqrt{1 - 1/n} \cdot l/\sqrt{n}$, the minimal volume of any such parallelepiped is $(1 - \epsilon)^n l^n$. In particular the volume is nonzero.⁷ Since

$$\epsilon = \frac{\sqrt{n}}{2}S < \sqrt{n}S\sqrt{1 - \frac{1}{n}} = \sqrt{1 - \frac{1}{n}} \cdot \frac{l}{\sqrt{n}},$$

the volume of $\mathcal{P}(\text{Rot}(\mathbf{c}))$ is nonzero, and the matrix $\text{Rot}(\mathbf{c})$ has full rank. \square

In [48], Lemma 2.12 is used to prove that the expected squared norm $\|\mathbf{d} - \mathbf{c}\|^2$ (when \mathbf{d} is chosen according to distribution $D_{\Lambda, s, \mathbf{c}}$) is at most $s^2 \cdot n$. In this paper we will need a bound on the expected value of the convolution product $\|(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x}\|^2$. It immediately follows from the result in [48] and inequality (2.6) that for any vector \mathbf{x} , the expectation of $\|(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x}\|^2$ is at most $s^2 \cdot n^2 \cdot \|\mathbf{x}\|^2$. Below, we use Lemma 2.12 to directly prove a stronger bound.

Lemma 3.2 *For any n -dimensional lattice Λ , positive reals $\epsilon \leq 1/3$, $s \geq 2\eta_\epsilon(\Lambda)$ and vectors $\mathbf{c}, \mathbf{x} \in \mathbb{R}^n$,*

$$E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\|(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x}\|^2] \leq s^2 \cdot n \cdot \|\mathbf{x}\|^2.$$

Proof: Let $\mathbf{e}_1, \dots, \mathbf{e}_n$ be the standard basis of \mathbb{R}^n . Notice that $(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x} = \mathbf{x} \otimes (\mathbf{d} - \mathbf{c}) = \text{Rot}(\mathbf{x}) \cdot (\mathbf{d} - \mathbf{c})$, and $\mathbf{e}_i^T \cdot \text{Rot}(\mathbf{x}) = (\text{rot}^i(\tilde{\mathbf{x}}))^T$, where $\tilde{\mathbf{x}} = (x_n, \dots, x_1)^T$ is the reverse of \mathbf{x} . By linearity of expectation, we have

$$E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\|(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x}\|^2] = \sum_{i=1}^n E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\langle \mathbf{e}_i, (\mathbf{d} - \mathbf{c}) \otimes \mathbf{x} \rangle^2].$$

For every $i = 1, \dots, n$,

$$\begin{aligned} \langle \mathbf{e}_i, (\mathbf{d} - \mathbf{c}) \otimes \mathbf{x} \rangle &= \mathbf{e}_i^T \cdot \text{Rot}(\mathbf{x}) \cdot (\mathbf{d} - \mathbf{c}) \\ &= \langle \text{rot}^i(\tilde{\mathbf{x}}), \mathbf{d} - \mathbf{c} \rangle \\ &= \|\mathbf{x}\| \langle \mathbf{u}_i, \mathbf{d} - \mathbf{c} \rangle \end{aligned}$$

⁷The minimal volume $(1 - \epsilon)^n l^n$ is achieved by the intuitive solution that shortens each edge by ϵ . Interestingly, as shown in [45], when $\epsilon = l/\sqrt{n}$ there are better ways to choose the perturbations that result in a singular parallelepiped with zero volume.

where $\mathbf{u}_i = \text{rot}^i(\tilde{\mathbf{x}})/\|\mathbf{x}\|$ is a unit vector. So,

$$E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\|(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x}\|^2] = \|\mathbf{x}\|^2 \cdot \sum_{i=1}^n E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\langle \mathbf{u}_i, \mathbf{d} - \mathbf{c} \rangle^2].$$

Using the assumption $s \geq 2\eta_\epsilon(\Lambda)$ and applying Lemma 2.12, we get that for all $i = 1, \dots, n$,

$$E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\langle \mathbf{u}_i, \mathbf{d} - \mathbf{c} \rangle^2] \leq s^2 \left(\frac{1}{2\pi} + \frac{\epsilon}{1 - \epsilon} \right) \leq s^2 \left(\frac{1}{2\pi} + \frac{1/3}{1 - 1/3} \right) \leq s^2.$$

Adding up for all i and substituting in the previous equation we get

$$E_{\mathbf{d} \sim D_{\Lambda, s, \mathbf{c}}} [\|(\mathbf{d} - \mathbf{c}) \otimes \mathbf{x}\|^2] \leq s^2 \|\mathbf{x}\|^2 n.$$

□

4 Generalized compact knapsacks

The hash function families considered in this paper, as well as in previous works [2, 10, 46, 48], are all special cases of the following general definition.

Definition 4.1 For any ring R , subset $S \subset R$ and integer $m \geq 1$, the generalized knapsack function family $\mathcal{H}(R, S, m) = \{f_{\mathbf{a}} : S^m \rightarrow R\}_{\mathbf{a} \in R^m}$ is defined by

$$f_{\mathbf{a}}(\mathbf{x}) = \sum_{i=1}^m x_i \cdot a_i,$$

for all $\mathbf{a} \in R^m$ and $\mathbf{x} \in S^m$, where $\sum_i x_i \cdot a_i$ is computed using the ring addition and multiplication operations.

In this paper we consider the ring $R = (\mathbb{F}_{p(n)}^n, +, \otimes)$ of n -dimensional vectors over the finite field $\mathbb{F}_{p(n)}$ with $p(n) = n^{\Theta(1)}$ elements, with the usual vector addition operation and convolution product \otimes . For brevity, we will denote this ring simply as $\mathbb{F}_{p(n)}^n$. We remark that for any prime p , the field \mathbb{F}_p is isomorphic to the ring \mathbb{Z}_p of integers modulo p . Here we use notation \mathbb{F}_p^n instead of \mathbb{Z}_p^n both because some of our results are valid even when p is not a prime, and also to emphasize that \mathbb{F}_p^n is the ring of vectors with the convolution product operation, rather than the componentwise multiplication of the product ring \mathbb{Z}_p^n .

As for S , we consider the set $S = D^n \subset \mathbb{F}_p^n$ of vectors with entries in an appropriately selected subset of \mathbb{F}_p . We want to study the hash function family $\mathcal{H}(\mathbb{F}_p^n, D^n, m)$, and prove that it is both almost regular and one-way.

The rest of the section is organized as follows. In Subsection 4.1 we prove that $\mathcal{H}(\mathbb{F}_p^n, D^n, m)$ is almost regular. In Subsection 4.2 we introduce and start studying a new worst-case lattice problem that will be instrumental to prove our main result. In Subsection 4.3 we give a reduction from solving this problem in the worst case to the problem of inverting functions $\mathcal{H}(\mathbb{F}_p^n, D^n, m)$ on the average. Finally, in Subsection 4.4, we use reductions among worst-case problems to establish the hardness of inverting $\mathcal{H}(\mathbb{F}_p^n, D^n, m)$ on the average based on the worst-case intractability of various standard problems (like SIVP and GDD) on cyclic lattices.

4.1 Regularity lemma

For any ring R of size $|R| \geq 2^n$, a necessary condition for the hash function family $\mathcal{H}(R, \{0, 1\}, m)$ to be almost regular is $m \geq \Omega(\log |R|) \geq \Omega(n)$, because when $m \leq o(\log |R|)$, at most a tiny fraction of the elements of R can be expressed as the sum of a subset of $\{a_1, \dots, a_m\}$. In this subsection we prove that the hash function family $\mathcal{H}(\mathbb{F}_p^n, D^n, m)$ is almost regular already when $m = \omega(1)$ is an unbounded function with

arbitrarily slow growth rate. Our proof is quite different from the standard proof for the subset-sum function $\mathcal{H}(R, \{0, 1\}, m)$. In particular, while the proof for $\mathcal{H}(R, \{0, 1\}, m)$ only relies on the additive structure of R , our proof makes full use of the ring properties of \mathbb{F}_p^n and the characterization of its ideals as quotients of polynomial rings.

Theorem 4.2 *For any finite field \mathbb{F} , subset $D \subset \mathbb{F}$, and integers n, m , the hash function family $\mathcal{H}(\mathbb{F}^n, D^n, m)$ is ϵ -regular for*

$$\epsilon = \frac{1}{2} \sqrt{(1 + |\mathbb{F}|/|D|^m)^n - 1}.$$

In particular, for any $p(n) = n^{O(1)}$, $|D| = n^{\Omega(1)}$ and $m(n) = \omega(1)$, the function ensemble $\mathcal{H}(\mathbb{F}_{p(n)}^n, D^n, m(n))$ is almost regular.

The proof of the theorem is based on the following lemma attributed to Rackoff by Impagliazzo and Zuckerman.

Lemma 4.3 ([28, Claim 2]) *Let V, V' be independent and identically distributed random variables taking values in a finite set S . If V, V' have collision probability $\Pr\{V = V'\} \leq (1 + 4\epsilon^2)/|S|$, then the statistical distance between V and the uniform distribution over S is at most ϵ .*

Proof: For completeness, we give a sketch of the proof. Using the second inequality in (2.4), the statistical distance between V and $U(S)$ can be bounded by

$$\frac{1}{2} \sum_{s \in S} \left| \Pr\{V = s\} - \frac{1}{|S|} \right| \leq \frac{1}{2} \sqrt{|S|} \sqrt{\sum_{s \in S} (\Pr\{V = s\} - 1/|S|)^2}.$$

Expanding the square and adding up for all $s \in S$, the last expression can be rewritten as

$$\frac{1}{2} \sqrt{|S|} \sqrt{\Pr\{V = V'\} - \frac{2}{|S|} + \frac{1}{|S|}} = \frac{1}{2} \sqrt{|S| \cdot \Pr\{V = V'\} - 1}$$

which is at most ϵ under the assumption that $\Pr\{V = V'\} \leq (1 + 4\epsilon^2)/|S|$. \square

We also need the following simple lemma.

Lemma 4.4 *Let R be a finite ring, and $z_1, \dots, z_m \in R$ a sequence of arbitrary ring elements. If $a_1, \dots, a_m \in R$ are independently and uniformly distributed ring elements, then $\sum a_i \cdot z_i$ is uniformly distributed over the ideal $\langle z_1, \dots, z_m \rangle$ generated by z_1, \dots, z_m . In particular, for any $z_1, \dots, z_m \in R$ and randomly chosen $a_1, \dots, a_m \in R$, the probability that $\sum a_i \cdot z_i = 0$ is exactly $1/|\langle z_1, \dots, z_m \rangle|$.*

Proof: Let $z_1, \dots, z_m \in R$ be arbitrary ring elements, and, for any $b \in R$, define $A_b = \{(a_1, \dots, a_m) \in R^m : \sum a_i \cdot z_i = b\}$. Notice that the probability that $\sum a_i \cdot z_i = b$ (over the random choice of a_1, \dots, a_m) equals $|A_b|/|R|^m$. If $b \notin \langle z_1, \dots, z_m \rangle$, then $A_b = \emptyset$ and $\Pr\{\sum a_i \cdot z_i = b\} = 0$. It remains to prove that all $b \in \langle z_1, \dots, z_m \rangle$ have the same probability. Let $b = \sum a_i \cdot z_i$ be an arbitrary element of $\langle z_1, \dots, z_m \rangle$. We claim that $|A_b| = |A_0|$. It is easy to see that $\mathbf{a}' \in A_b$ if and only if $\mathbf{a}' - \mathbf{a} \in A_0$. Since $\mathbf{a}' \mapsto \mathbf{a}' - \mathbf{a}$ is a bijection between A_b and A_0 , it follows that $|A_b| = |A_0|$. This proves that all $b \in \langle z_1, \dots, z_m \rangle$ have the same probability $|A_b|/|R|^m = |A_0|/|R|^m$, and completes the proof of the lemma. \square

We are now ready to prove the theorem.

Proof [of Theorem 4.2]: We want to prove that $\text{OWF}(\mathcal{H}(\mathbb{F}^n, D^n, m))$ is very close to the uniform distribution over $(\mathbb{F}^n)^m \times \mathbb{F}^n$. We first bound the collision probability of two independent copies of $\text{OWF}(\mathcal{H}(\mathbb{F}^n, D^n, m))$. Let $((\mathbf{a}_1, \dots, \mathbf{a}_m), \sum_i \mathbf{a}_i \otimes \mathbf{x}_i)$ and $((\mathbf{a}'_1, \dots, \mathbf{a}'_m), \sum_i \mathbf{a}'_i \otimes \mathbf{x}'_i)$ be two independent samples chosen according to

the distribution $\text{OWF}(\mathcal{H}(\mathbb{F}^n, D^n, m))$. By definition, the elements $\mathbf{a}_i, \mathbf{a}'_i \in \mathbb{F}^n$ and $\mathbf{x}_i, \mathbf{x}'_i \in D^n$ are all chosen independently and uniformly at random from their respective sets. Therefore, the collision probability is

$$\begin{aligned} \Pr \left\{ \forall i. \mathbf{a}_i = \mathbf{a}'_i \wedge \sum_{i=1}^m \mathbf{a}_i \otimes \mathbf{x}_i = \sum_{i=1}^m \mathbf{a}'_i \otimes \mathbf{x}'_i \right\} &= \Pr \{ \forall i. \mathbf{a}_i = \mathbf{a}'_i \} \\ &\cdot \Pr \left\{ \sum_{i=1}^m \mathbf{a}_i \otimes \mathbf{x}_i = \sum_{i=1}^m \mathbf{a}'_i \otimes \mathbf{x}'_i \mid \forall i. \mathbf{a}_i = \mathbf{a}'_i \right\} \\ &= \frac{1}{|\mathbb{F}|^{mn}} \cdot \Pr \left\{ \sum_{i=1}^m \mathbf{a}_i \otimes (\mathbf{x}_i - \mathbf{x}'_i) = \mathbf{0} \right\}. \end{aligned}$$

By Lemma 4.4, the probability (over the random choice of $\mathbf{a}_1, \dots, \mathbf{a}_m$) that $\sum_i \mathbf{a}_i \otimes (\mathbf{x}_i - \mathbf{x}'_i) = \mathbf{0}$ equals $1/|I|$ where $I = \langle \mathbf{x}_1 - \mathbf{x}'_1, \dots, \mathbf{x}_m - \mathbf{x}'_m \rangle$ is the ideal generated by $\mathbf{x}_1 - \mathbf{x}'_1, \dots, \mathbf{x}_m - \mathbf{x}'_m$. Let \mathcal{I} be the set of all ideals of $(\mathbb{F}^n, +, \otimes)$. Conditioning on the ideal I , the collision probability can be expressed as

$$\begin{aligned} \frac{1}{|\mathbb{F}|^{mn}} \cdot \Pr \left\{ \sum_{i=1}^m \mathbf{a}_i \otimes (\mathbf{x}_i - \mathbf{x}'_i) = \mathbf{0} \right\} &= \frac{1}{|\mathbb{F}|^{nm}} \cdot \sum_{I \in \mathcal{I}} \frac{\Pr \{ \langle \mathbf{x}_1 - \mathbf{x}'_1, \dots, \mathbf{x}_m - \mathbf{x}'_m \rangle = I \}}{|I|} \\ &\leq \frac{1}{|\mathbb{F}|^{nm}} \cdot \sum_{I \in \mathcal{I}} \frac{\Pr \{ \langle \mathbf{x}_1 - \mathbf{x}'_1, \dots, \mathbf{x}_m - \mathbf{x}'_m \rangle \subseteq I \}}{|I|} \\ &= \frac{1}{|\mathbb{F}|^{n(m+1)}} \cdot \sum_{I \in \mathcal{I}} \frac{|\mathbb{F}|^n}{|I|} \cdot \prod_{i=1}^m \Pr \{ (\mathbf{x}_i - \mathbf{x}'_i) \in I \}. \end{aligned}$$

In the rest of the proof, we regard \mathbb{F}^n as the ring of univariate polynomials $\mathbb{F}[\alpha]$ modulo $\alpha^n - 1$. Since \mathbb{F} is a field, $\mathbb{F}[\alpha]$ is a principal ideal domain, i.e., all ideals in $\mathbb{F}[\alpha]$ are of the form $\langle Q(\alpha) \rangle$ for some polynomial $Q(\alpha) \in \mathbb{F}[\alpha]$. It follows that all ideals $I \in \mathcal{I}$ of the quotient ring $\mathbb{F}[\alpha]/(\alpha^n - 1)$ are of the form $\langle Q(\alpha) \rangle$ where $Q(\alpha)$ is a factor of $\alpha^n - 1$. (To see this, given an ideal $I \in \mathcal{I}$, select a representative for each element of I , and let $Q(\alpha)$ be the greatest common divisor of all these representatives and the polynomial $\alpha^n - 1$.) Let $(\alpha^n - 1) = Q_1(\alpha) \cdot Q_2(\alpha) \cdot \dots \cdot Q_r(\alpha)$ be the factorization of $(\alpha^n - 1)$ into irreducible polynomials over \mathbb{F} , and for any subset $S \subseteq \{1, \dots, r\}$, let $Q_S(\alpha) = \prod_{i \in S} Q_i(\alpha)$. The ideals of R are $\mathcal{I} = \{ \langle Q_S \rangle : S \subseteq \{1, \dots, r\} \}$. For any ideal $\langle Q_S \rangle \in \mathcal{I}$, we have $|\langle Q_S \rangle| = |\mathbb{F}|^{n - \deg(Q_S)}$ and

$$\Pr \{ (\mathbf{x}_i - \mathbf{x}'_i) \in \langle Q_S \rangle \} = \Pr \{ \mathbf{x}_i \equiv \mathbf{x}'_i \pmod{Q_S} \} \leq \max_{\tilde{\mathbf{x}}} \Pr \{ \mathbf{x}_i \pmod{Q_S} = \tilde{\mathbf{x}} \} \leq \frac{1}{|D|^{\deg(Q_S)}}, \quad (4.8)$$

where $\tilde{\mathbf{x}}$ ranges over all polynomials of degree $\deg(\tilde{\mathbf{x}}) < \deg(Q_S)$, and the last inequality follows from the fact that, for any fixed value of the $(n - \deg(Q_S))$ higher order coefficients of \mathbf{x} , the function $\mathbf{x} \mapsto \mathbf{x} \pmod{Q_S}$ is a bijection between sets of size $|D|^{\deg(Q_S)}$. Using the bound (4.8), we get

$$\frac{|\mathbb{F}|^n}{|\langle Q_S \rangle|} \prod_{i=1}^m \Pr \{ (\mathbf{x}_i - \mathbf{x}'_i) \in \langle Q_S \rangle \} \leq \frac{|\mathbb{F}|^n}{|\mathbb{F}|^{n - \deg(Q_S)}} \left(\frac{1}{|D|^{\deg(Q_S)}} \right)^m = \left(\frac{|\mathbb{F}|}{|D|^m} \right)^{\deg(Q_S)}$$

and, adding up over all ideals,

$$\begin{aligned} \sum_{(Q_S) \in \mathcal{I}} \frac{|\mathbb{F}|^n}{|\langle Q_S \rangle|} \prod_{i=1}^m \Pr \{ (\mathbf{x}_i - \mathbf{x}'_i) \in \langle Q_S \rangle \} &\leq \sum_S \left(\frac{|\mathbb{F}|}{|D|^m} \right)^{\deg(Q_S)} \\ &= \prod_{i=1}^r \left(1 + \left(\frac{|\mathbb{F}|}{|D|^m} \right)^{\deg(Q_i)} \right) \\ &\leq \prod_{i=1}^r \left(1 + \frac{|\mathbb{F}|}{|D|^m} \right)^{\deg(Q_i)} \\ &= \left(1 + \frac{|\mathbb{F}|}{|D|^m} \right)^n. \end{aligned}$$

This proves that the collision probability is at most

$$\frac{(1 + |\mathbb{F}|/|D|^m)^n}{|\mathbb{F}|^{n(m+1)}}.$$

Now observe that random variable $\text{OWF}(\mathcal{H}(\mathbb{F}^n, D^n, m))$ takes values in the set $(\mathbb{F}^n)^m \times \mathbb{F}^n$, which has size $|\mathbb{F}|^{n(m+1)}$. Therefore, by Lemma 4.3, the statistical distance between $\text{OWF}(\mathcal{H}(\mathbb{F}^n, D^n, m))$ and the uniform distribution over $(\mathbb{F}^n)^m \times \mathbb{F}^n$ is at most

$$\epsilon = \frac{1}{2} \sqrt{\left(1 + \frac{|\mathbb{F}|}{|D|^m}\right)^n - 1}.$$

□

4.2 The worst case problems

We want to show that inverting our generalized compact knapsack function $\mathcal{H}(\mathbb{F}^n, D^n, m)$ (on the average and with non-negligible probability) is at least as hard as solving GDD_γ (as well as various other related problems) over cyclic lattices in the worst case. Following [46], this is done in two steps. First, all relevant worst-case lattice problems are reduced to an intermediate worst-case problem, and then the intermediate problem is reduced to the problem of inverting functions in $\mathcal{H}(\mathbb{F}^n, D^n, m)$ on the average. In [46], the goal was to reduce the worst-case problem SIVP_γ to the problem of inverting⁸ $\mathcal{H}(\mathbb{Z}_p^n, \{0, 1\}, m)$ on the average, and the intermediate problem was an incremental version of SIVP , where given a lattice basis \mathbf{B} , a set of sufficiently long linearly independent lattice vectors \mathbf{S} , and a hyperplane H , the goal is to find a lattice vector not in H shorter than $\|\mathbf{S}\|$ by some constant factor.

Here we consider a different intermediate problem, which is an incremental version of GDD , where one is given a GDD instance (\mathbf{B}, \mathbf{t}) , a set of n linearly independent vectors $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$, and a sufficiently large real parameter r , and the goal is to find a lattice vector whose distance from the target is at most $\|\mathbf{S}\|/c + r$ for some constant c .

Definition 4.5 *The incremental guaranteed distance decoding problem ($\text{INCGDD}_{\gamma, c}^\phi$), given an n -dimensional lattice \mathbf{B} , a set of n linearly independent vectors $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$, a target $\mathbf{t} \in \mathbb{R}^n$, and a real $r > \gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$, asks for a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{s} - \mathbf{t}\| \leq (\|\mathbf{S}\|/c) + r$.*

In the rest of this subsection we show that many standard lattice problems reduce (via *lattice-preserving* reductions) to INCGDD . A reduction (say, from SIVP to INCGDD) is *lattice-preserving* if all calls to the INCGDD oracle are of the form $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ where \mathbf{B} is the SIVP input lattice. Lattice-preserving reductions are particularly useful in the context of our paper because they allow to reduce a (worst-case) lattice problem over a given class of lattices (e.g., cyclic lattices) to another (worst-case) lattice problem over the *same* class of lattices. For example, the next lemma shows that SIVP on cyclic lattices can be solved in polynomial time given oracle access to a procedure that solves INCGDD on cyclic lattices.

Lemma 4.6 *For any $\gamma(n) \geq 1$ and any ϕ , there exists a lattice-preserving reduction from $\text{SIVP}_{4\gamma}^\phi$ to $\text{INCGDD}_{\gamma, 5}^\phi$.*

Proof: Given a basis \mathbf{B} , our goal is to construct a set of n linearly independent vectors \mathbf{S} of length $\|\mathbf{S}\| \leq 4\gamma(n)\phi(\mathbf{B})$. We do this by an iterative process. Initially, we set \mathbf{S} to the result of applying the LLL basis reduction algorithm [35] to \mathbf{B} , so that \mathbf{S} is a basis for $\mathcal{L}(\mathbf{B})$ satisfying $\|\mathbf{S}\| \leq 2^n \lambda_n(\mathcal{L}(\mathbf{B}))$. At each step, we identify the longest vector in \mathbf{S} , say \mathbf{s}_i . We then take \mathbf{t} to be a vector orthogonal to $\mathbf{s}_1, \dots, \mathbf{s}_{i-1}, \mathbf{s}_{i+1}, \dots, \mathbf{s}_n$ of length $\|\mathbf{S}\|/2$. We call the INCGDD oracle with the instance $(\mathbf{B}, \mathbf{S}, \mathbf{t}, \|\mathbf{S}\|/4)$.

⁸In fact, [46] only requires an algorithm that finds collisions $f_{\mathbf{a}}(\mathbf{x}) = f_{\mathbf{a}}(\mathbf{x}')$, an easier problem than inverting the function $f_{\mathbf{a}}$.

If it fails, we terminate and output \mathbf{S} . Otherwise, we obtain a lattice vector \mathbf{s} within distance at most $(\|\mathbf{S}\|/5) + \|\mathbf{S}\|/4 = (9/20)\|\mathbf{S}\|$ from \mathbf{t} . Notice that $\|\mathbf{s}\| \leq \|\mathbf{t}\| + \|\mathbf{s} - \mathbf{t}\| \leq (19/20)\|\mathbf{S}\|$. Moreover, \mathbf{s} is linearly independent from $\{\mathbf{s}_1, \dots, \mathbf{s}_{i-1}, \mathbf{s}_{i+1}, \dots, \mathbf{s}_n\}$ because it is at distance $\|\mathbf{t}\| - \|\mathbf{s} - \mathbf{t}\| \geq \|\mathbf{S}\|/20 > 0$ from the hyperplane spanned by those vectors. So, we can replace \mathbf{s}_i with \mathbf{s} and repeat the process.

Notice that when the oracle call fails, it must be the case that $\|\mathbf{S}\|/4 \leq \gamma(n)\phi(\mathbf{B})$, and hence $\|\mathbf{S}\| \leq 4\gamma(n)\phi(\mathbf{B})$, as required. It remains to bound the running time of the reduction. Every iteration takes time polynomial in the size of \mathbf{B} and \mathbf{S} , which is polynomial in the size of the original input \mathbf{B} because \mathbf{S} is initially equal to a polynomial time computable function of \mathbf{B} and $\|\mathbf{S}\|$ can only get smaller from one iteration to the next. Finally, consider the quantity $\log \prod_i \|\mathbf{s}_i\|$. We know that initially $\|\mathbf{s}_i\| \leq 2^n \lambda_n(\mathcal{L}(\mathbf{B}))$ for all i . Moreover, $\log \prod_i \|\mathbf{s}_i\|$ decreases by a constant (namely, $\log 19/20$) at each iteration, and it is always at least as large as $\log \prod_i \min\{\|\mathbf{s}_i\|, (10/11)\lambda_n(\mathcal{L}(\mathbf{B}))\}$. (To see this, observe that the vector selected for replacement at every iteration must satisfy $\|\mathbf{s}_i\| = \max_i \|\mathbf{s}_i\| \geq \lambda_n(\mathcal{L}(\mathbf{B}))$.) Therefore, the procedure terminates after at most $O(\log \prod_i (11/10)2^n) = O(n^2)$ iterations. \square

Next, we show how to use an INCGDD oracle to solve GDD.

Lemma 4.7 *For any $\gamma(n) \geq 1$ and any ϕ , there exists a lattice-preserving reduction from $\text{GDD}_{2\gamma}^\phi$ to $\text{INCGDD}_{\gamma,5}^\phi$.*

Proof: Given a basis \mathbf{B} and a vector \mathbf{t} , our goal is to find a lattice vector within distance $2\gamma(n)\phi(\mathbf{B})$ of \mathbf{t} . First, we apply the reduction in Lemma 4.6 to obtain a set $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$ of n linearly independent vectors of length at most $\|\mathbf{S}\| \leq 4\gamma(n)\phi(\mathbf{B})$. Then, we apply the LLL basis reduction algorithm [35] to \mathbf{B} to obtain a basis for the same lattice such that $\|\mathbf{B}'\| \leq 2^n \lambda_n(\mathbf{B})$. Define $r = (\sqrt{n}/4)\|\mathbf{B}'\|$ and call the INCGDD oracle on input $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$. If the oracle returns an invalid answer,⁹ it must be that $r \leq \gamma(n)\phi(\mathbf{B})$, and we can efficiently find a lattice vector \mathbf{s} within distance $(\sqrt{n}/2)\|\mathbf{B}'\| = 2r \leq 2\gamma(n)\phi(\mathbf{B})$ from the target \mathbf{t} using Babai's nearest plane algorithm [7]. So, assume the oracle call $\text{INCGDD}(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ succeeds. We keep calling the INCGDD oracle with smaller and smaller values of r , until either a call returns an invalid answer or r gets too small. Specifically, at each iteration we decrease r by a factor $10/11$, as long as $r > \|\mathbf{S}\|/4$. Notice that if $r \leq \|\mathbf{S}\|/4$ and the oracle $\text{INCGDD}(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ returns a valid answer \mathbf{s} , then we can terminate with output \mathbf{s} because

$$\|\mathbf{s} - \mathbf{t}\| \leq r + \|\mathbf{S}\|/5 \leq (1/4 + 1/5)\|\mathbf{S}\| < \|\mathbf{S}\|/2 \leq 2\gamma(n)\phi(\mathbf{B}).$$

Finally, assume that the oracle succeeds for some value of r , but fails in the next iteration. Since the oracle fails on input $(10/11)r$, it must be that $(10/11)r \leq \gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$, or, equivalently, $r \leq (11/10)\gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$. Therefore, the vector \mathbf{s} returned by the oracle on input r satisfies

$$\text{dist}(\mathbf{s}, \mathbf{t}) \leq r + \frac{\|\mathbf{S}\|}{5} \leq \frac{11}{10}\gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B})) + \frac{4}{5}\gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B})) < 2\gamma(n) \cdot \phi(\mathcal{L}(\mathbf{B}))$$

as required.

This concludes the description of the reduction. The reduction is correct because in every case, it terminates with a lattice vector within distance $2\gamma(n)\phi(\mathcal{L}(\mathbf{B}))$ from the target. Moreover, it is clear that each iteration can be implemented in time polynomial in the size of the original \mathbf{B} . In order to complete the proof we only need to make sure that the number of iterations is also polynomial. Notice that since $\|\mathbf{S}\| \geq \lambda_n(\mathbf{B})$ (by definition of λ_n) and $\|\mathbf{B}'\| \leq 2^n \lambda_n(\mathbf{B})$ (by the properties of LLL reduced basis), the number of iterations is at most $\log_{11/10}(\sqrt{n}\|\mathbf{B}'\|/\|\mathbf{S}\|) \leq \log_{11/10}(\sqrt{n}2^n) = O(n)$, which is polynomial in the size of \mathbf{B} . \square

4.3 The main reduction

In this section we reduce the worst-case problem $\text{INCGDD}_{\gamma,c}^\eta$ on cyclic lattices to the problem of inverting the generalized compact knapsack functions $\mathcal{H}(\mathbb{F}_{p(n)}^n, D^n, m(n))$ on the average.

⁹We remark that the validity of the answer \mathbf{s} returned by the oracle can be easily checked by verifying that $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ and $\|\mathbf{s} - \mathbf{t}\| \leq r + \|\mathbf{S}\|/5$.

Our reduction is closely related to Micciancio and Regev’s variant [48] of Ajtai’s original worst-case average-case connection [2] for arbitrary lattices. The average-case problem considered in [2, 48] is that of finding a small integer linear dependency among randomly chosen elements of the additive group \mathbb{Z}_p^n . This problem is shown to be at least as hard as finding short vectors in an arbitrary lattice \mathbf{B} roughly¹⁰ as follows:

1. First, the group \mathbb{Z}_p^n is embedded into the space spanned by the lattice \mathbf{B} by dividing each side of the fundamental parallelepiped $\mathcal{P}(\mathbf{B})$ into p equal parts, yielding p^n *subregions* naturally corresponding to the group elements.
2. Next, samples from the group \mathbb{Z}_p^n are chosen by selecting (almost) uniformly distributed vectors $\mathbf{y}'_i \in \mathcal{P}(\mathbf{B})$, and rounding them to the corners \mathbf{a}_i of the corresponding *subregions* of $\mathcal{P}(\mathbf{B})$. The fundamental idea of [48] (also used in this paper) is that random points $\mathbf{y}'_i \in \mathcal{P}(\mathbf{B})$ can be efficiently generated together with nearby lattice vectors by choosing a *relatively short* random vector \mathbf{y}_i (with Gaussian distribution), and reducing it modulo the lattice. This yields a point $\mathbf{y}'_i = \mathbf{y}_i \bmod \mathbf{B}$ distributed almost uniformly at random within $\mathcal{P}(\mathbf{B})$, together with a nearby lattice point $(\mathbf{y}'_i - \mathbf{y}_i) \in \mathcal{L}(\mathbf{B})$.
3. Since the vectors \mathbf{y}'_i are (almost) uniformly distributed over $\mathcal{P}(\mathbf{B})$, the associated rounded vectors \mathbf{a}_i correspond to (almost) uniformly chosen group elements in \mathbb{Z}_p^n . So, one can use the average-case oracle to find a small integer linear dependency relation (x_1, \dots, x_m) modulo p among the group elements. It is easy to see that $\sum_i \mathbf{a}_i \cdot x_i \in \mathcal{L}(\mathbf{B})$ for any such linear dependency relation (x_1, \dots, x_m) . Moreover, approximating each rounded vector \mathbf{a}_i with nearby lattice point $(\mathbf{y}'_i - \mathbf{y}_i)$, yields another lattice vector $\sum_i (\mathbf{y}'_i - \mathbf{y}_i) \cdot x_i \in \mathcal{L}(\mathbf{B})$ not far from $\sum_i \mathbf{a}_i \cdot x_i$. So, one can compute a short vector in the original lattice $\mathcal{L}(\mathbf{B})$ by taking the difference $\sum_i (\mathbf{a}_i - \mathbf{y}'_i + \mathbf{y}_i) \cdot x_i$ between these two lattice points.

The main novelty in our reduction is that we need to embed (into the space spanned by the lattice), not just the additive group \mathbb{Z}_p^n , but rather the more sophisticated ring structure $(\mathbb{Z}_p^n, +, \otimes)$ of n -dimensional vectors with the usual modular addition and new convolution product operation. Remarkably, our reduction achieves this by exploiting the rotational symmetry of cyclic lattices. Technically, the above reduction is modified by replacing the parallelepiped $\mathcal{P}(\mathbf{B})$ (used to embed \mathbb{Z}_p^n) with the region $\mathcal{P}(\mathbf{C})$ associated to an appropriately chosen sublattice. Namely, $\mathbf{C} = \text{Rot}(\mathbf{c})$ is a matrix obtained by taking the cyclic rotations of a single vector \mathbf{c} . The reason why this small modification makes the reduction work is mostly algebraic, rather than purely geometric as in [2, 48]. Still, our embedding inherits some of the geometric intuition of Ajtai’s proof [2], making the Gaussian techniques of Micciancio and Regev [48] applicable.

Another difference between the reductions of [2, 48] and ours, is that (for technical reasons) here we need to work with inhomogeneous variants of both the average-case and worst-case problems. For example, instead of considering the problem of finding a linear dependency among ring elements (i.e., a linear combination that adds up to zero), we consider the problem of inverting the generalized knapsack function (i.e., finding a linear combination of the knapsack weights that adds up to a given random target value). Again, the reason why this change is necessary to make our reduction work is mostly algebraic, not geometric.¹¹ So, while the geometric intuition behind Ajtai’s proof (as described above) can be useful to achieve a superficial understanding of our proof as well, we warn the reader that part of the proof is algebraic in nature and it is best understood by putting the geometric interpretation aside.

Theorem 4.8 *For any constants $c > 1$ and $\delta > 0$, negligible function $\epsilon(n) = n^{-\omega(1)}$, polynomially bounded integers $m(n) = n^{O(1)}$ and primes $p(n) = n^{O(1)}$, if $m(n) = \omega(1)$ and $p(n) \geq (3c \cdot m(n) \cdot n^{2.5})^{1/(1-\delta)}$, then there is a probabilistic polynomial time reduction from solving $\text{INCGDD}_{\gamma(n),c}^{\eta_\epsilon}$ within a factor $\gamma(n) = 3 \cdot m(n) \cdot n \cdot p(n)^\delta$ in the worst case over cyclic lattices (with high probability), to inverting the generalized compact knapsack function $\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n))$ on the average (with non-negligible probability).*

¹⁰Many important technical details are omitted in our informal description. The reader is referred to the original paper [48] for a more accurate description.

¹¹It has been recently shown [38, 57] that the problem of adapting our reduction to the homogeneous setting is intimately related to the factorization of $(\alpha^n - 1)$ into irreducible polynomials over $\mathbb{Z}[\alpha]$. The reader is referred to [38, 57] for details.

Proof: For any instance $\mathbf{Q} = [\mathbf{q}_1, \dots, \mathbf{q}_{m(n)}; \mathbf{q}_0] \in \mathbb{F}_{p(n)}^{n \cdot m(n)} \times \mathbb{F}_{p(n)}^n$ of the generalized compact knapsack function inversion problem, let

$$\Gamma(\mathbf{Q}) = \left\{ \mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{m(n)}]: \forall i. \mathbf{x}_i \in [p(n)^\delta]^n \wedge \sum_{i=1}^{m(n)} \mathbf{q}_i \otimes \mathbf{x}_i = \mathbf{q}_0 \pmod{p(n)} \right\}$$

be the corresponding set of solutions. Let \mathcal{F} be an oracle that on input an instance \mathbf{Q} selected at random according to distribution $\text{OWF}(\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n)))$, outputs a solution $\mathcal{F}(\mathbf{Q}) \in \Gamma(\mathbf{Q})$ with non-negligible probability. Let $\beta(n)$ be the probability that $\mathcal{F}(\tilde{\mathbf{Q}}) \in \Gamma(\tilde{\mathbf{Q}})$ when $\tilde{\mathbf{Q}}$ is selected uniformly at random from $\mathbb{F}_{p(n)}^{n \cdot m(n)} \times \mathbb{F}_{p(n)}^n$. Since $p(n) = n^{\Theta(1)}$, $[p(n)^\delta] \geq p(n)^\delta = n^{\Omega(1)}$ and $m(n) = \omega(1)$, by Theorem 4.2 the probability distribution $\text{OWF}(\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n)))$ is statistically close to the uniform one $U(\mathbb{F}_{p(n)}^{n \cdot m(n)} \times \mathbb{F}_{p(n)}^n)$. Therefore, $\beta(n)$ is non-negligible too. We use \mathcal{F} to solve problem $\text{INCGDD}_{\gamma, c}^\eta$ over cyclic lattices in the worst case, with non-negligible probability $\Omega(\beta(n))$. Since we are solving $\text{INCGDD}_{\gamma, c}^\eta$ in the worst case, the success probability of the reduction can be made exponentially close to 1 using standard repetition techniques.

Let $(\mathbf{B}, \mathbf{S}, \mathbf{t}, r)$ be a valid $\text{INCGDD}_{\gamma, c}^\eta$ instance such that the lattice $\mathcal{L}(\mathbf{B})$ is cyclic. We know that $\mathcal{L}(\mathbf{S})$ is a (not necessarily cyclic) full rank sublattice of $\mathcal{L}(\mathbf{B})$, and $r > \gamma(n) \cdot \eta_{\epsilon(n)}(\mathcal{L}(\mathbf{B}))$ for some negligible function $\epsilon(n) = n^{-\omega(1)}$. The goal of the reduction is to find a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ within distance $r + \|\mathbf{S}\|/c$ from the target \mathbf{t} . The reduction works as follows:

1. Use Lemma 3.1 to find a vector $\mathbf{c} \in \mathcal{L}(\mathbf{S}) \subseteq \mathcal{L}(\mathbf{B})$ of length $\|\mathbf{c}\|_1 \leq 2 \cdot n \cdot \|\mathbf{S}\|$ such that $\text{Rot}(\mathbf{c})$ has full rank.
2. For $i = 0, \dots, m(n)$, do the following
 - (a) Use Lemma 2.7 to generate a uniformly random lattice vector $\mathbf{v}_i \in \mathcal{L}(\mathbf{B}) \cap \mathcal{P}(\text{Rot}(\mathbf{c}))$:
 - (b) Generate a random noise vector \mathbf{y}_i with probability density $\mathbf{y}_i \sim \rho_s/s^n$ for $s = 2r/\gamma(n)$, and let $\mathbf{y}'_i = \mathbf{y}_i \pmod{\mathbf{B}}$.
 - (c) Compute $\mathbf{a}_i = \lfloor p(n) \cdot \text{Rot}(\mathbf{c})^{-1}(\mathbf{v}_i + \mathbf{y}'_i) \rfloor$.
3. Compute $\mathbf{b} = \lfloor p(n) \cdot \text{Rot}(\mathbf{c})^{-1} \mathbf{t} \rfloor$.
4. Define the generalized compact knapsack problem instance

$$\mathbf{Q} = [\mathbf{a}_1 \pmod{p(n)}, \dots, \mathbf{a}_{m(n)} \pmod{p(n)}; \mathbf{a}_0 + \mathbf{b} \pmod{p(n)}] \quad (4.9)$$

and invoke $\mathcal{F}(\mathbf{Q})$ to find a potential solution $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{m(n)}] \in \Gamma(\mathbf{Q})$.

5. Let $\mathbf{x}_0 = -\mathbf{e}_1$, and return the vector

$$\mathbf{s} = \sum_{i=0}^{m(n)} \left(\mathbf{v}_i + \mathbf{y}'_i - \frac{\mathbf{c} \otimes \mathbf{a}_i}{p(n)} - \mathbf{y}_i \right) \otimes \mathbf{x}_i + \frac{\mathbf{c} \otimes \mathbf{b}}{p(n)}.$$

The correctness of the reduction is based on the following two lemmas. The first lemma shows that if the oracle \mathcal{F} successfully outputs a solution $\mathbf{X} \in \Gamma(\mathbf{Q})$, then the reduction outputs a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$.

Lemma 4.9 *If $\mathbf{X} \in \Gamma(\mathbf{Q})$, then $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ is a lattice vector.*

Proof: Let $\mathbf{X} \in \Gamma(\mathbf{Q})$ be a valid solution to knapsack instance (4.9). In particular, assume¹² $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_{m(n)})$ satisfies

$$\sum_{i=1}^{m(n)} \mathbf{a}_i \otimes \mathbf{x}_i \equiv (\mathbf{a}_0 + \mathbf{b}) \pmod{p(n)}. \quad (4.10)$$

¹²In fact, this is the only property of \mathbf{X} required in this lemma. The other property $\forall i. \mathbf{x}_i \in [p(n)^\delta]^n$ is not used in the proof.

Using the distributive and associative properties of \otimes , vector \mathbf{s} can be rewritten as the sum

$$\mathbf{s} = \sum_{i=0}^{m(n)} (\mathbf{v}_i + \mathbf{y}'_i - \mathbf{y}_i) \otimes \mathbf{x}_i - \mathbf{c} \otimes \frac{\sum_{i=0}^{m(n)} \mathbf{a}_i \otimes \mathbf{x}_i - \mathbf{b}}{p(n)}.$$

We claim that all terms in the summation belong to the lattice $\mathcal{L}(\mathbf{B})$. First of all notice that for any $i \geq 0$, the vector $\mathbf{v}_i + \mathbf{y}'_i - \mathbf{y}_i$ belongs to the lattice $\mathcal{L}(\mathbf{B})$ because $\mathbf{v}_i \in \mathcal{L}(\mathbf{B})$ and $\mathbf{y}'_i \equiv \mathbf{y}_i$ modulo $\mathcal{L}(\mathbf{B})$. Using the cyclicity of $\mathcal{L}(\mathbf{B})$, we get that all columns of $\text{Rot}(\mathbf{v}_i + \mathbf{y}'_i - \mathbf{y}_i)$ belong to the lattice, and

$$(\mathbf{v}_i + \mathbf{y}'_i - \mathbf{y}_i) \otimes \mathbf{x}_i = \text{Rot}(\mathbf{v}_i + \mathbf{y}'_i - \mathbf{y}_i) \cdot \mathbf{x}_i \in \mathcal{L}(\mathbf{B})$$

because \mathbf{x}_i has integer entries. For the last term, we use equalities (4.10) and $\mathbf{a}_0 \otimes \mathbf{x}_0 = -\mathbf{a}_0$, yielding

$$\sum_{i \geq 0} \mathbf{a}_i \otimes \mathbf{x}_i - \mathbf{b} = \sum_{i \geq 1} \mathbf{a}_i \otimes \mathbf{x}_i - (\mathbf{a}_0 + \mathbf{b}) \equiv \mathbf{0} \pmod{p(n)}.$$

Therefore $(\sum_{i \geq 0} \mathbf{a}_i \otimes \mathbf{x}_i - \mathbf{b})/p(n)$ is an integer vector, and

$$\mathbf{c} \otimes \frac{\sum_{i \geq 0} \mathbf{a}_i \otimes \mathbf{x}_i - \mathbf{b}}{p(n)} = \text{Rot}(\mathbf{c}) \cdot \frac{\sum_{i \geq 0} \mathbf{a}_i \otimes \mathbf{x}_i - \mathbf{b}}{p(n)} \in \mathcal{L}(\text{Rot}(\mathbf{c})) \subseteq \mathcal{L}(\mathbf{B})$$

belongs to the lattice. \square

The second lemma shows that the input \mathbf{Q} to the oracle \mathcal{F} is almost uniformly distributed, and therefore $\mathcal{F}(\mathbf{Q})$ is successful with probability very close to $\beta(n)$.

Lemma 4.10 *For any $s \geq \eta_{\epsilon(n)}(\mathcal{L}(\mathbf{B}))$, the statistical distance of \mathbf{Q} (as defined in (4.9)) from the uniform distribution is at most*

$$\frac{1}{2}(m(n) + 1) \cdot \epsilon(n).$$

In particular, for any polynomially bounded $m(n) = n^{O(1)}$, and negligible function $\epsilon(n) = n^{-\omega(1)}$, the distribution of \mathbf{Q} is within negligible distance from the uniform distribution $U(\mathbb{F}_{p(n)}^{n \cdot m(n)} \times \mathbb{F}_{p(n)}^n)$.

Proof: We first bound the distance of each $\mathbf{a}_i \pmod{p(n)}$ from the uniform distribution over $\mathbb{F}_{p(n)}^n$. Notice that

$$\begin{aligned} \mathbf{a}_i \pmod{p(n)} &= \lfloor p(n) \cdot \text{Rot}(\mathbf{c})^{-1}(\mathbf{v}_i + \mathbf{y}'_i) \rfloor \pmod{p(n)} \\ &= \lfloor p(n) \cdot \text{Rot}(\mathbf{c})^{-1}((\mathbf{v}_i + \mathbf{y}'_i) \pmod{\text{Rot}(\mathbf{c})}) \rfloor. \end{aligned}$$

So, if \mathbf{y}'_i were distributed uniformly at random over $\mathcal{P}(\mathbf{B})$, then $(\mathbf{v}_i + \mathbf{y}'_i) \pmod{\text{Rot}(\mathbf{c})}$ would be uniform over $\mathcal{P}(\text{Rot}(\mathbf{c}))$, and $\mathbf{a}_i \pmod{p(n)}$ would also have perfectly uniform distribution over $\mathbb{F}_{p(n)}^n$. Consider $\mathbf{a}_i \pmod{p(n)}$ as a randomized function of \mathbf{y}'_i , i.e., define the randomized function $g(\mathbf{y}') = \lfloor p(n) \cdot \text{Rot}(\mathbf{c})^{-1}(\mathbf{v} + \mathbf{y}') \rfloor \pmod{p(n)}$ where $\mathbf{v} \in \mathcal{L}(\mathbf{B}) \cap \mathcal{P}(\text{Rot}(\mathbf{c}))$ is chosen uniformly at random. Notice that $g(\mathbf{y}'_i) = \mathbf{a}_i \pmod{p(n)}$. We observed that g maps the uniform distribution over $\mathcal{P}(\mathbf{B})$ to the uniform distribution over $\mathbb{F}_{p(n)}^n$, i.e., $g(U(\mathcal{P}(\mathbf{B}))) = U(\mathbb{F}_{p(n)}^n)$. Therefore, by Proposition 2.2 the statistical distance between $\mathbf{a}_i \pmod{p(n)}$ and the uniform distribution over $\mathbb{F}_{p(n)}^n$ is at most

$$\begin{aligned} \Delta(\mathbf{a}_i \pmod{p(n)}, U(\mathbb{F}_{p(n)}^n)) &= \Delta(g(\mathbf{y}'_i), g(U(\mathcal{P}(\mathbf{B})))) \\ &\leq \Delta(\mathbf{y}'_i, U(\mathcal{P}(\mathbf{B}))). \end{aligned}$$

Notice that \mathbf{y}'_i has distribution $\rho_s/s^n \pmod{\mathcal{P}(\mathbf{B})}$. Using the assumption $s \geq \eta_{\epsilon}(\mathcal{L}(\mathbf{B}))$ and Lemma 2.11, we get that

$$\Delta(\mathbf{a}_i \pmod{p(n)}, U(\mathbb{F}_{p(n)}^n)) \leq \Delta(\mathbf{y}'_i, U(\mathcal{P}(\mathbf{B}))) \leq \epsilon(n)/2.$$

Now consider the knapsack instance \mathbf{Q} defined in (4.9). Since the elements of \mathbf{Q} are independently distributed, by Proposition 2.4 we have

$$\Delta(\mathbf{Q}, U(\mathbb{F}_{p(n)}^{n \cdot m(n)} \times \mathbb{F}_{p(n)}^n)) \leq \sum_{i=1}^{m(n)} \Delta(\mathbf{a}_i \bmod p(n), U(\mathbb{F}_{p(n)}^n)) + \Delta(\mathbf{a}_0 + \mathbf{b} \bmod p(n), U(\mathbb{F}_{p(n)}^n)).$$

The last term satisfies

$$\Delta(\mathbf{a}_0 + \mathbf{b} \bmod p(n), U(\mathbb{F}_{p(n)}^n)) = \Delta(\mathbf{a}_0 \bmod p(n), (U(\mathbb{F}_{p(n)}^n) - \mathbf{b}) \bmod p(n)) = \Delta(\mathbf{a}_0 \bmod p(n), U(\mathbb{F}_{p(n)}^n)).$$

Therefore,

$$\Delta(\mathbf{Q}, U(\mathbb{F}_{p(n)}^{n \cdot m(n)} \times \mathbb{F}_{p(n)}^n)) \leq \sum_{i=0}^{m(n)} \Delta(\mathbf{a}_i \bmod p(n), U(\mathbb{F}_{p(n)}^n)) \leq (m(n) + 1) \cdot \epsilon(n)/2.$$

□

We are now ready to prove the correctness of the reduction. Namely, we want to prove that for any n -dimensional cyclic lattice basis \mathbf{B} , full rank subset $\mathbf{S} \subset \mathcal{L}(\mathbf{B})$, target \mathbf{t} , and $r > \gamma(n) \cdot \eta_\epsilon(\mathcal{L}(\mathbf{B}))$, the reduction outputs a lattice vector $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ such that $\|\mathbf{s} - \mathbf{t}\| \leq r + \|\mathbf{S}\|/c$ with non-negligible probability $\Omega(\beta(n))$. By Lemma 4.9, $\mathbf{s} \in \mathcal{L}(\mathbf{B})$ is satisfied whenever oracle \mathcal{F} returns a valid solution $\mathbf{X} = \mathcal{F}(\mathbf{Q}) \in \Gamma(\mathbf{Q})$. Therefore, the success probability of the reduction is at least

$$\begin{aligned} \Pr \left\{ \mathbf{s} \in \mathcal{L}(\mathbf{B}), \|\mathbf{s} - \mathbf{t}\| \leq r + \frac{\|\mathbf{S}\|}{c} \right\} &\geq \Pr \left\{ \mathbf{X} \in \Gamma(\mathbf{Q}), \|\mathbf{s} - \mathbf{t}\| \leq r + \frac{\|\mathbf{S}\|}{c} \right\} \\ &= \Pr \{ \mathbf{X} \in \Gamma(\mathbf{Q}) \} \cdot \Pr \left\{ \|\mathbf{s} - \mathbf{t}\| \leq r + \frac{\|\mathbf{S}\|}{c} \mid \mathbf{X} \in \Gamma(\mathbf{Q}) \right\} \end{aligned} \quad (4.11)$$

Let $\tilde{\mathbf{Q}} \in U(\mathbb{F}_{p(n)}^{n \times m(n)} \times \mathbb{F}_{p(n)}^n)$ be an instance distributed uniformly at random. Notice that $s = 2r/\gamma(n) > 2\eta_\epsilon(n)(\mathcal{L}(\mathbf{B})) > \eta_\epsilon(n)(\mathcal{L}(\mathbf{B}))$. So, by Lemma 4.10, $\Delta(\mathbf{Q}, \tilde{\mathbf{Q}})$ is negligible. Therefore, the first probability in (4.11) satisfies

$$\begin{aligned} \Pr \{ \mathbf{X} \in \Gamma(\mathbf{Q}) \} &= \Pr \{ \mathcal{F}(\mathbf{Q}) \in \Gamma(\mathbf{Q}) \} \\ &\geq \Pr \{ \mathcal{F}(\tilde{\mathbf{Q}}) \in \Gamma(\tilde{\mathbf{Q}}) \} - \Delta(\mathbf{Q}, \tilde{\mathbf{Q}}) \\ &= \beta(n) - n^{-\omega(1)} \geq \Omega(\beta(n)). \end{aligned}$$

We bound the second probability in (4.11) using Markov's inequality:

$$\begin{aligned} \Pr \left\{ \|\mathbf{s} - \mathbf{t}\| \leq r + \frac{\|\mathbf{S}\|}{c} \mid \mathbf{X} \in \Gamma(\mathbf{Q}) \right\} &= 1 - \Pr \left\{ \|\mathbf{s} - \mathbf{t}\| > r + \frac{\|\mathbf{S}\|}{c} \mid \mathbf{X} \in \Gamma(\mathbf{Q}) \right\} \\ &\geq 1 - \frac{E \left[\|\mathbf{s} - \mathbf{t}\| \mid \mathbf{X} \in \Gamma(\mathbf{Q}) \right]}{r + \|\mathbf{S}\|/c}. \end{aligned} \quad (4.12)$$

We will prove that the conditional expectation $E \left[\|\mathbf{s} - \mathbf{t}\| \mid \mathbf{X} \in \Gamma(\mathbf{Q}) \right]$ is at most $(2/3)(1 + 2/m(n)) \cdot (r + \|\mathbf{S}\|/c)$, so that (4.12) is at least

$$1 - \frac{2}{3} \left(1 + \frac{2}{m(n)} \right) = \frac{1}{3} - \frac{4}{3m(n)} = \Omega(1).$$

This proves that (4.11) (and therefore also the success probability of the reduction) is at least $\Omega(\beta(n)) \cdot \Omega(1) = \Omega(\beta(n))$. It remains to bound the expected length of $\mathbf{s} - \mathbf{t}$. By the triangle inequality,

$$\|\mathbf{s} - \mathbf{t}\| \leq \sum_{i=0}^{m(n)} \left\| \left(\mathbf{v}_i + \mathbf{y}'_i - \frac{\mathbf{c} \otimes \mathbf{a}_i}{p(n)} \right) \otimes \mathbf{x}_i \right\| + \sum_{i=0}^{m(n)} \|\mathbf{y}_i \otimes \mathbf{x}_i\| + \left\| \frac{\mathbf{c} \otimes \mathbf{b}}{p(n)} - \mathbf{t} \right\|. \quad (4.13)$$

We will show that the first and third term can be made arbitrarily small by using a sufficiently large value for $p(n)$. (Here is where we use the assumption $p(n) \geq (3cm(n)n^{2.5})^{1/(1-\delta)}$.) The second term requires the most effort and it is bounded using the properties of Gaussian distributions. Notice that

$$\mathbf{v}_i + \mathbf{y}'_i - \frac{\mathbf{c} \otimes \mathbf{a}_i}{p(n)} = \frac{\mathbf{c} \otimes (\mathbf{w} - \lfloor \mathbf{w} \rfloor)}{p(n)}$$

where $\mathbf{w} = p(n) \text{Rot}(\mathbf{c})^{-1}(\mathbf{v}_i + \mathbf{y}'_i)$. Since $\|\mathbf{c}\|_1 \leq 2n\|\mathbf{S}\|$ by construction and $\|\mathbf{w} - \lfloor \mathbf{w} \rfloor\|_\infty \leq 1$ for any vector \mathbf{w} , we have

$$\left\| \mathbf{v}_i + \mathbf{y}'_i - \frac{\mathbf{c} \otimes \mathbf{a}_i}{p(n)} \right\|_\infty \leq \frac{\|\mathbf{c}\|_1 \cdot \|\mathbf{w} - \lfloor \mathbf{w} \rfloor\|_\infty}{p(n)} \leq \frac{2n\|\mathbf{S}\|}{p(n)}. \quad (4.14)$$

Similarly, we have

$$\left\| \mathbf{t} - \frac{\mathbf{c} \otimes \mathbf{b}}{p(n)} \right\|_\infty \leq \frac{2n\|\mathbf{S}\|}{p(n)}. \quad (4.15)$$

Multiplying (4.14) by \mathbf{x}_i and using $\|\mathbf{x}_i\|_1 \leq n \cdot p(n)^\delta$, we get

$$\left\| \left(\mathbf{v}_i + \mathbf{y}'_i - \frac{\mathbf{c} \otimes \mathbf{a}_i}{p(n)} \right) \otimes \mathbf{x}_i \right\|_\infty \leq \left\| \mathbf{v}_i + \mathbf{y}'_i - \frac{\mathbf{c} \otimes \mathbf{a}_i}{p(n)} \right\|_\infty \cdot \|\mathbf{x}_i\|_1 \leq \frac{2n^2\|\mathbf{S}\|}{p(n)^{1-\delta}}. \quad (4.16)$$

Substituting (4.15) and (4.16) in (4.13), and using the second inequality in (2.5)

$$\begin{aligned} \|\mathbf{s} - \mathbf{t}\| &\leq (m(n) + 1) \cdot \frac{2n^{2.5}\|\mathbf{S}\|}{p(n)^{1-\delta}} + \frac{2n^{1.5}\|\mathbf{S}\|}{p(n)} + \sum_{i=0}^{m(n)} \|\mathbf{y}_i \otimes \mathbf{x}_i\| \\ &\leq 2(m(n) + 2) \cdot \frac{n^{2.5}\|\mathbf{S}\|}{p(n)^{1-\delta}} + \sum_{i=0}^{m(n)} \|\mathbf{y}_i \otimes \mathbf{x}_i\|. \end{aligned}$$

Using the assumption $p(n) \geq (3c \cdot m(n) \cdot n^{2.5})^{1/(1-\delta)}$, the first term in the last expression is at most

$$2(m(n) + 2) \cdot \frac{n^{2.5}\|\mathbf{S}\|}{p(n)^{1-\delta}} \leq \frac{2}{3} \left(1 + \frac{2}{m(n)} \right) \cdot \frac{\|\mathbf{S}\|}{c}.$$

We want to prove that the conditional expectation of the second term satisfies

$$E \left[\sum_{i=0}^{m(n)} \|\mathbf{y}_i \otimes \mathbf{x}_i\| \mid \mathbf{X} \in \Gamma(\mathbf{Q}) \right] \leq \frac{2}{3} \left(1 + \frac{2}{m(n)} \right) \cdot r, \quad (4.17)$$

so that the expectation of $\|\mathbf{s} - \mathbf{t}\|$ is at most $(2/3)(1 + 2/m(n))(r + \|\mathbf{S}\|/c)$ as claimed.

In order to prove (4.17) we fix the value of \mathbf{X} , \mathbf{Q} and $\mathbf{Y}' = [\mathbf{y}'_0, \dots, \mathbf{y}'_{m(n)}]$, and consider the conditional expectation of $\|\mathbf{s} - \mathbf{t}\|$ given all these values. We will show that for any *fixed* \mathbf{Q} , \mathbf{X} and \mathbf{Y}' (satisfying $\mathbf{X} \in \Gamma(\mathbf{Q})$), the conditional expectation of $\|\mathbf{s} - \mathbf{t}\|$ satisfies the bound in (4.17). Equation (4.17) immediately follows by averaging over all possible values of \mathbf{Q} , \mathbf{X} and \mathbf{Y}' such that $\mathbf{X} \in \Gamma(\mathbf{Q})$. So, in the rest of the proof, we fix the value of \mathbf{Q} , \mathbf{X} and \mathbf{Y}' , and consider the conditional distribution of the vectors \mathbf{y}_i . Notice that, given \mathbf{y}'_i , vector \mathbf{y}_i must necessarily belong to the set $\mathbf{y}'_i + \mathcal{L}(\mathbf{B})$, but it is otherwise random and independent from \mathbf{Q} , \mathbf{X} , \mathbf{Y}' and all other \mathbf{y}_j 's. So, the conditional distribution of $\mathbf{y}_i \in \mathbf{y}'_i + \mathcal{L}(\mathbf{B})$ is

$$\Pr \{ \mathbf{y}_i \mid \mathbf{Y}', \mathbf{Q}, \mathbf{X} \} = \Pr \{ \mathbf{y}_i \mid \mathbf{y}'_i \} = \frac{\rho_s(\mathbf{y}_i)}{\rho_s(\mathbf{y}'_i + \mathcal{L}(\mathbf{B}))} = \frac{\rho_{s, -\mathbf{y}'_i}(\mathbf{y}_i - \mathbf{y}'_i)}{\rho_{s, -\mathbf{y}'_i}(\mathcal{L}(\mathbf{B}))}.$$

In other words, the conditional distribution of $(\mathbf{y}_i - \mathbf{y}'_i) \in \mathcal{L}(\mathbf{B})$ is $D_{\mathcal{L}(\mathbf{B}), s, -\mathbf{y}'_i}$. Recall that $s = 2r/\gamma(n) > 2\eta_{\epsilon(n)}(\mathcal{L}(\mathbf{B}))$. So, by Lemma 3.2,

$$\begin{aligned} E \left[\|\mathbf{y}_i \otimes \mathbf{x}_i\|^2 \mid \mathbf{y}'_i \right] &= E_{(\mathbf{y}_i - \mathbf{y}'_i) \sim D_{\mathcal{L}(\mathbf{B}), s, -\mathbf{y}'_i}} \left[\|((\mathbf{y}_i - \mathbf{y}'_i) - (-\mathbf{y}'_i)) \otimes \mathbf{x}_i\|^2 \right] \\ &\leq s^2 \|\mathbf{x}_i\|^2 n \\ &\leq s^2 n^2 \cdot p(n)^{2\delta}. \end{aligned}$$

By convexity, we get

$$E [\| \mathbf{y}_i \otimes \mathbf{x}_i \| \mid \mathbf{y}'_i] \leq n \cdot s \cdot p(n)^\delta.$$

Finally, adding up for all values of i and using the definition of $s = 2r/\gamma(n)$ and $\gamma(n) = 3m(n) \cdot n \cdot p(n)^\delta$, we get

$$\begin{aligned} \sum_{i=0}^{m(n)} E [\| \mathbf{y}_i \otimes \mathbf{x}_i \| \mid \mathbf{y}'_i] &\leq (m(n) + 1) \cdot n \cdot s \cdot p(n)^\delta \\ &= \frac{2r(m(n) + 1) \cdot n \cdot p(n)^\delta}{\gamma(n)} \\ &= \left(1 + \frac{1}{m(n)} \right) \frac{2r}{3} \leq \frac{2}{3} \left(1 + \frac{2}{m(n)} \right) r. \end{aligned}$$

This concludes the proof that the conditional expectation $E [\| \mathbf{s} - \mathbf{t} \| \mid \mathbf{X} \in \Gamma(\mathbf{Q})]$ is at most $(2/3)(1 + 2/m(n))(r + \|\mathbf{S}\|/c)$, and the reduction succeeds with non-negligible probability $\Omega(\beta(n))$. \square

By choosing a small enough $\delta > 0$ in the previous theorem, we obtain the following corollary.

Corollary 4.11 *For every $\alpha > 0$ there is a $\delta > 0$ such that for every constant $c > 1$, negligible function $\epsilon(n) > 0$, primes $p(n) = \Theta(n^3)$, and subpolynomial integers $m(n) = n^{o(1)}$ satisfying $m(n) = \omega(1)$, there is a probabilistic polynomial time reduction from solving $\text{INCGDD}_{\gamma,c}^{\eta,\epsilon}$ in the worst case within a factor $\gamma(n) = n^{1+\alpha}$ to inverting $\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n))$ on the average with non-negligible probability.*

Proof: Assume without loss of generality that $\alpha < 1/2$, and take for example $\delta = \alpha/6$. Notice that $p(n)^{1-\delta} \geq \Theta(n^{3-\alpha/2})$ is asymptotically bigger than $3cm(n)n^{2.5} \leq O(n^{3-\alpha})$. Therefore, for all sufficiently large n , $p(n)$ satisfies the hypothesis of Theorem 4.8, and inverting $\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n))$ on the average is at least as hard as solving $\text{INCGDD}_{\gamma,c}^\eta$ in the worst case, for

$$\gamma(n) = 3m(n)np(n)^\delta \leq n^{1+o(1)+\alpha/2} \leq n^{1+\alpha}.$$

\square

4.4 Other lattice problems

In Subsection 4.3 we have shown that inverting the generalized compact knapsack functions $\mathcal{H}(\mathbb{F}_p^n, [p^\delta]^n, \omega(1))$ on the average is at least as hard as solving the INCGDD problem over cyclic lattices in the worst case. In this subsection we relate the complexity of inverting the compact knapsack functions to well known worst-case lattice problems restricted to cyclic lattices.

Corollary 4.12 *For any $\alpha > 0$ there is a $\delta > 0$ such that for all primes $p(n) = \Theta(n^3)$ and integers $m(n) = n^{o(1)}$ satisfying $m(n) = \omega(1)$, inverting $\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n))$ on the average with non-negligible probability is at least as hard as solving any of the following problems in the worst case within a factor $\gamma(n) = n^{1+\alpha}$:*

- the guaranteed distance decoding problem GDD_γ^η over cyclic lattices
- the generalized independent vector problem SIVP_γ^η over cyclic lattices.

Proof: Both reductions easily follow by combining Corollary 4.11 with Lemmas 4.6 and 4.7. \square

Finally, using known relations between η and λ_n (see Lemma 2.13) and $\lambda_n \leq 2\rho$ (see [47, Theorem 7.9]), we can relate the hardness of breaking one-way function $\mathcal{H}(\mathbb{F}_p^n, [p^\delta]^n, \omega(1))$ to the standard version of the lattice problems GDD and SIVP.

Corollary 4.13 *For any $\alpha > 0$ there is a $\delta > 0$ such that for any primes $p(n) = \Theta(n^3)$ and integers $m(n) = n^{o(1)}$ satisfying $m(n) = \omega(1)$, inverting $\mathcal{H}(\mathbb{F}_{p(n)}^n, [p(n)^\delta]^n, m(n))$ on the average with non-negligible probability is at least as hard as solving any of the following problems in the worst case for $\gamma(n) = n^{1+\alpha}$:*

- *the guaranteed distance decoding problem GDD_γ over cyclic lattices*
- *the shortest independent vector problem SIVP_γ over cyclic lattices.*

5 Remarks and open problems

We have introduced a new class of very efficient one-way functions with strong security guarantees. Namely, our functions are provably hard to invert (on the average), based on a worst-case intractability assumption. The assumption is that no polynomial time algorithm can approximate SIVP, GDD, or other related lattice problems, in the worst case over cyclic lattices within a factor $n^{1+\epsilon}$ almost linear in the dimension of the lattice.

Similarly to the case of general lattices [2, 10, 19, 46, 48], our results too can be interpreted as a connection between the worst-case and average-case complexity of various lattice problems. In [2, 10, 19, 46, 48] it is shown that finding small nonzero integer solutions to a random linear equation $\mathbf{Ax} = \mathbf{0} \pmod p$ on the average is at least as hard as solving SIVP and other lattice problems in the worst case. Since the integer solutions to the equation

$$\Lambda(\mathbf{A}) = \{\mathbf{x}: \mathbf{Ax} = \mathbf{0} \pmod p\}$$

form a lattice, the results in [2, 10, 19, 46, 48] can be formulated as a reduction from solving SIVP in the worst case to solving SVP on the average.

In this paper we have shown that inverting our generalized compact knapsack functions on the average is at least as hard as the worst case instance of GDD, as well as other lattices problems, over cyclic lattices. We now show how inverting the compact knapsack function can also be formulated as a lattice problem. A compact knapsack function $\mathbf{a}_1, \dots, \mathbf{a}_m$ implicitly defines a lattice in dimension $O(m \cdot n)$ given by the set of all $(\mathbf{y}_1^T, \dots, \mathbf{y}_m^T)^T$ such that $\sum \mathbf{a}_i \otimes \mathbf{y}_i = \mathbf{0}$. In fact, using matrix notation, one can consider the weights $\mathbf{a}_1, \dots, \mathbf{a}_m$ as a compact representation of an $n \times m \cdot n$ matrix

$$\mathbf{A} = [\text{Rot}(\mathbf{a}_1) | \dots | \text{Rot}(\mathbf{a}_m)]$$

which defines a lattice $\Lambda(\mathbf{A}) = \{\mathbf{x}: \mathbf{Ax} = \mathbf{0} \pmod p\}$ in the usual way. Up to a permutation of the coordinates, it is immediate to see that the lattice associated to matrix \mathbf{A} above is quasi-cyclic of order m , i.e., it is invariant under shifts rot^m by m positions. Inverting the subset-sum function can be formulated as a closest vector problem instance as follows. Given $\mathbf{a}_1, \dots, \mathbf{a}_m$, and knapsack target \mathbf{b} , we first compute an arbitrary solution $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_m)$ to the equation $\sum \mathbf{a}_i \otimes \mathbf{z}_i = \mathbf{b}$. (These vectors \mathbf{z}_i are not required to belong to $S = D^n$, and can be efficiently found.) Then, finding small vectors $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_m)$ such that $\sum \mathbf{a}_i \otimes \mathbf{x}_i = \mathbf{b}$ is equivalent to finding lattice vectors $(\mathbf{z}_1 - \mathbf{x}_1, \dots, \mathbf{z}_m - \mathbf{x}_m) \in \Lambda(\mathbf{A})$ close to $(\mathbf{z}_1, \dots, \mathbf{z}_m)$.

So, our result can be interpreted as follows: if GDD on n -dimensional cyclic lattices is hard to approximate within $n^{1+\epsilon}$ factors in the *worst case*, then CVP on $\omega(n)$ dimensional $\omega(1)$ -cyclic lattices is hard to solve on the *average*.

Many open problems remain. First and foremost, since the security of the cryptographic function proposed in this paper relies on the hardness of lattice problems on cyclic lattices, it would be very interesting to investigate the worst case complexity of computational problems on cyclic lattices, as discussed in the introduction. In the rest of this section we describe other open problems (and recent results) concerning the construction of various cryptographic primitives and improvement of the worst-case inapproximability factor required by the proof of security.

Cryptographic applications. From a practical point of view, it would be nice to prove that our function satisfies stronger security guarantees than one-wayness. For the case of general lattices, it is known [19, 48]

that under the assumption that SIVP is hard to approximate in the worst case within almost linear factors $\omega(n \log n)$, the generalized subset-sum function over \mathbb{Z}_p^n is not only one-way, but also collision resistant. Unfortunately, technical differences between our proof and the one in [48] make it hard to establish the same result for the compact knapsack function. At the time this paper was written, we left as an open problem to prove or disprove that our generalized compact knapsack function is collision resistant. This problem has been independently settled in [38] and [57], where it is shown that the generalized compact knapsack function proposed in this paper is not collision resistant (in fact, it is not even a universal one-way hash function), but it can be turned into a collision resistant hash function (under essentially the same worst-case complexity assumption) by suitably restricting its domain or modifying the underlying ring. In [57, 38] it is also shown that the function is secure even when the number of weights $m(n) = O(1)$ is constant, improving over the almost constant bound $m(n) = \omega(1)$ used in this paper.

It would be interesting to use the techniques developed in this paper to build (very efficient) cryptographic primitives other than one-way functions and collision resistant hash functions. One-way functions are known to be sufficient to build many other useful cryptographic primitives, like pseudo-random generators [21, 24], one-way hash functions [50], commitment schemes [49], digital signatures schemes [61], or private key encryption schemes [20]. However, these generic constructions are rather inefficient, so with their use most of the efficiency benefits of our compact knapsack function would be lost. We leave as an open problem the construction of any such provably secure cryptographic primitive with efficiency comparable to our one-way function, and based on similar worst-case intractability assumptions. As already mentioned above, for the case of hash functions, this problem has been recently solved in [38, 57], where it is shown that a simple variant of the function proposed in this paper is collision resistant, and therefore also a one-way hash function. We also observe that by standard reductions [50, 14], the collision resistant hash function of [57, 38] also yields efficient constructions of (statistically hiding) commitment schemes based on the worst-case hardness of cyclic lattices.

Finally, and probably the hardest of the open problems concerning the cryptographic applicability of our techniques, is to build a public-key encryption scheme (or a trapdoor function) with efficiency and security guarantees similar to our compact knapsack function. Building *public-key* encryption schemes seems a much harder problem than building one-way functions or private key encryptions. Still, we believe that designing public-key encryption schemes with efficiency and security properties similar to our one-way function may not be so out of reach. We remark that the class of cyclic lattices used in this paper is related to (although different from) the class of “convolutional modular lattices” used by NTRU [25], a commercial public-key cryptosystem based on lattices. Specifically, the lattices used by NTRU can be described as quasi-cyclic lattices of order 2, i.e., lattices that are invariant under cyclic shifts by 2 positions. Unfortunately, no proof of security is known for NTRU (even based on nontrivial *average-case* complexity assumptions). Still, based on the similarities between NTRU and other lattice-based cryptosystems [42], we hope that, as Ajtai’s one-way function [2] inspired the design of public-key cryptosystems [3, 59], our work will provide a starting point for the design of *efficient* and provably secure public-key cryptosystems based on *cyclic* lattices. Proving the security of NTRU, or finding alternative ways to build public-key cryptosystems with efficiency and security properties similar to our one-way function is left as an open problem.

Improving the connection factor. The worst-case inapproximability factor for SIVP and GDD required by our one-way function is $n^{1+\epsilon}$, for arbitrarily small $\epsilon > 0$. A modest improvement has been recently achieved in [57, 38], who proved that (a variant of) the generalized compact knapsack proposed in this paper is as hard to invert on the average (in fact, even collision resistant) as approximating SIVP (on cyclic lattices) within a factor $n \log^{O(1)} n$, essentially matching similar results for general lattices [48]. An interesting open question is whether it is possible to do even better than that. We remark that the worst-case problems solved by our reduction are somehow harder than SIVP and GDD. Our reduction allows to solve SIVP $^\eta$ and GDD $^\eta$ within almost linear factors, and then uses known relations between the smoothing parameter η and standard lattice parameters like λ_n and ρ . An interesting question is whether better relations between η , λ_n and ρ can be proved in the case of cyclic lattices.

For the case of GDD, we showed how to solve GDD $^{\lambda_n}$ within almost linear factors $n^{1+\epsilon}$, and then used

the inequality $\rho \geq \lambda_n/2$ to express our result in terms of $\text{GDD} = \text{GDD}^\rho$. Since ρ can be larger than λ_n by $\sqrt{n}/2$ (even for the case of cyclic lattices), our reduction may approximate GDD within factors much smaller than $n^{1+\epsilon}$, potentially as low as $n^{0.5+\epsilon}$, depending on the input lattice. We leave as an open problem to prove that the generalized compact knapsack function is as hard to invert as approximating GDD over cyclic lattices in the worst case within factors $\gamma(n) = n^{0.5+\epsilon}$.

6 Acknowledgments

The author would like to thank Oded Regev, Vadim Lyubashevsky and the anonymous referees for useful discussions and comments.

References

- [1] M. Ajtai. The shortest vector problem in l_2 is NP-hard for randomized reductions (extended abstract). In *Proceedings of the thirtieth annual ACM symposium on theory of computing - STOC '98*, pages 10–19, Dallas, Texas, USA, May 1998. ACM.
- [2] M. Ajtai. Generating hard instances of lattice problems. *Complexity of Computations and Proofs, Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.
- [3] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the twenty-ninth annual ACM symposium on theory of computing - STOC '97*, pages 284–293, El Paso, Texas, USA, May 1997. ACM.
- [4] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of the thirty-third annual ACM symposium on theory of computing - STOC 2001*, pages 266–275, Heraklion, Crete, Greece, July 2001. ACM.
- [5] H. Amirazizi, E. Karnin, and J. Reyneri. Compact knapsacks are polynomially solvable. *ACM SIGACT News*, 15:20–22, 1983. Preliminary version in Crypto 1981.
- [6] S. Arora, L. Babai, J. Stern, and E. Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, Apr. 1997. Preliminary version in FOCS'93.
- [7] L. Babai. On Lovasz' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [8] J. Blömer and J.-P. Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *Proceedings of the thirty-first annual ACM symposium on theory of computing - STOC '99*, pages 711–720, Atlanta, Georgia, USA, May 1999. ACM.
- [9] E. F. Brickell. Breaking iterated knapsacks. In G. R. Blakley and D. Chaum, editors, *Advances in cryptology - Proceedings of CRYPTO '84*, volume 196 of *Lecture Notes in Computer Science*, pages 342–358, Santa Barbara, California, USA, Aug. 1984. Springer.
- [10] J.-Y. Cai and A. P. Nerurkar. An improved worst-case to average-case connection for lattice problems (extended abstract). In *Proceedings of the 38th annual symposium on foundations of computer science - FOCS '97*, pages 468–477, Miami Beach, Florida, USA, Oct. 1997. IEEE.
- [11] B. Chor and R. Rivest. A knapsack-type public key cryptosystem based on arithmetic in finite fields. *IEEE Transactions in Information Theory*, 34:901–909, 1988. Preliminary version in CRYPTO'84.

- [12] M. J. Coster, A. Joux, B. A. LaMacchia, A. M. Odlyzko, C.-P. Schnorr, and J. Stern. Improved low-density subset sum algorithms. *Computational Complexity*, 2(2):111–128, 1992. Preliminary versions in Eurocrypt '91 and FCT '91.
- [13] T. W. Cusick. Cryptanalysis of a public key system based on Diophantine equations. *Information Processing Letters*, 56(2):73–75, 27 Oct. 1995.
- [14] I. Damgård, T. P. Pedersen, and B. Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. *Journal of Cryptology*, 10(3):163–194, 1997. Preliminary version in CRYPTO 1993.
- [15] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003. Preliminary version in FOCS 1998.
- [16] U. Feige and D. Micciancio. The inapproximability of lattice and coding problems with preprocessing. *Journal of Computer and System Sciences*, 69(1):45–67, 2003. Preliminary version in CCC 2002.
- [17] A. Flaxman and B. Przydatek. Solving medium-density subset sum problems in expected polynomial time. In V. Diekert and B. Durand, editors, *Proceedings of the 22nd annual symposium on Theoretical Aspects of computer science - STACS 2005*, volume 3404 of *Lecture Notes in Computer Science*, pages 305–314, Stuttgart, Germany, Jan. 2005. Springer.
- [18] C. Gentry and M. Szydło. Cryptanalysis of the revised NTRU signature scheme. In L. Knudsen, editor, *Advances in cryptology - EUROCRYPT 2002, Proceedings of the international conference on the theory and application of cryptographic techniques*, volume 2332 of *Lecture Notes in Computer Science*, pages 299–320, Amsterdam, The Netherlands, Apr. 2002. Springer.
- [19] O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. Technical Report TR96-056, Electronic Colloquium on Computational Complexity (ECCC), 1996.
- [20] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33:792–807, 1986.
- [21] O. Goldreich and L. Levin. A hard predicate for all one-way functions. In *Proceedings of the twenty-first annual ACM symposium on the theory of computing - STOC '89*, Seattle, Washington, USA, May 1989. ACM.
- [22] R. M. F. Goodman and A. J. McAuley. A new trapdoor knapsack public-key cryptosystem. In T. Beth, N. Cot, and I. Ingemarsson, editors, *Advances in cryptology - EUROCRYPT '84, Proceedings of a workshop on the theory and application of cryptographic techniques*, volume 209 of *Lecture Notes in Computer Science*, pages 150–158, Paris, France, Apr. 1984. Springer.
- [23] V. Guruswami and A. Vardy. Maximum-likelihood decoding of Reed-Solomon codes is NP-hard. *IEEE Transactions on Information Theory*, 51(7):2249–2256, July 2005. Preliminary version in SODA 2005.
- [24] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [25] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a ring based public key cryptosystem. In J. P. Buhler, editor, *Algorithmic number theory: Third international symposium - ANTS-III*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288, Portland, OR, USA, June 1998. Springer.
- [26] N. Howgrave-Graham and M. Szydło. A method to solve cyclotomic norm equations. In D. A. Buell, editor, *Algorithmic number theory: 6th international symposium - ANTS-VI*, volume 3076 of *Lecture Notes in Computer Science*, pages 272–279, Burlington, VT, USA, June 2004. Springer.

- [27] R. Impagliazzo and M. Naor. Efficient cryptographic schemes provably as secure as subset sum. *Journal of Cryptology*, 9(4):199–216, 1996.
- [28] R. Impagliazzo and D. Zuckerman. How to recycle random bits. In *Proceedings of the 30th annual symposium on foundations of computer science - FOCS '89*, pages 248–253, Research Triangle Park, NC, USA, Oct. 1989. IEEE.
- [29] A. Joux and J. Stern. Cryptanalysis of another knapsack cryptosystem. In H. Imai, R.L. Rivest, and T. Matsumoto, editors, *Advances in Cryptology – Proceedings Asiacrypt'91*, volume 739 of *Lecture Notes in Computer Science*, pages 470–476. Springer-Verlag, 1993.
- [30] A. Joux and J. Stern. Lattice reduction: A toolbox for the cryptanalyst. *Journal of Cryptology*, 11(3):161–185, 1998.
- [31] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of computer computation*, pages 85–103. Plenum, 1972.
- [32] S. Khot. Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM*, 52(5):789–808, Sept. 2005. Preliminary version in FOCS 2004.
- [33] J. C. Lagarias and A. M. Odlyzko. Solving low-density subset sum problems. *Journal of the ACM*, 32(1):229–246, Jan. 1985.
- [34] M.-K. Lee and K. Park. Low-density attack of public-key cryptosystems based on compact knapsacks. *Journal of Electrical Engineering and Information Science*, 4(2):197–204, 1999.
- [35] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:513–534, 1982.
- [36] C. H. Lin, C. C. Chang, and R. C. T. Lee. A new public-key cipher system based upon the Diophantine equations. *IEEE Transactions on Computers*, 44(1):13–19, 1995.
- [37] V. Lyubashevsky. The parity problem in the presence of noise, decoding random linear codes, and the subsetsum problem. In C. Chekuri, K. Jansen, D. P. J. Rolim, and L. Trevisan, editors, *APPROX-RANDOM 2005*, volume 3624 of *Lecture Notes in Computer Science*, pages 378–389, Berkeley, CA, USA, Aug. 2005. Springer.
- [38] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In I. Wegener, V. Sassone, and B. Preneel, editors, *Proceedings of the 33rd international colloquium on automata, languages and programming - ICALP 2006*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155, Venice, Italy, July 2006. Springer.
- [39] A. May and J. H. Silverman. Dimension reduction methods for convolution modular lattices. In J. Silverman, editor, *Cryptography and lattices conference – CaLC 2001*, volume 2146 of *Lecture Notes in Computer Science*, pages 110–125, Providence, RI, USA, Mar. 2001. Springer.
- [40] R. Merkle and M. Hellman. Hiding information and signatures in trapdoor knapsacks. *IEEE Transactions on Information Theory*, 24(5):525–530, Sept. 1978.
- [41] D. Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Transactions on Information Theory*, 47(3):1212–1215, Mar. 2001.
- [42] D. Micciancio. Improving lattice based cryptosystems using the Hermite normal form. In J. Silverman, editor, *Cryptography and lattices conference – CaLC 2001*, volume 2146 of *Lecture Notes in Computer Science*, pages 126–145, Providence, RI, USA, Mar. 2001. Springer.
- [43] D. Micciancio. The shortest vector problem is NP-hard to approximate to within some constant. *SIAM Journal on Computing*, 30(6):2008–2035, Mar. 2001. Preliminary version in FOCS 1998.

- [44] D. Micciancio. Generalized compact knapsaks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *Proceedings of the 43rd annual symposium on foundations of computer science - FOCS 2002*, pages 356–365, Vancouver, British Columbia, Canada, Nov. 2002. IEEE.
- [45] D. Micciancio. A note on the minimal volume of almost cubic parallelepiped. *Discrete and Computational Geometry*, 29(1):133–138, Dec. 2002.
- [46] D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai’s connection factor. *SIAM Journal on Computing*, 34(1):118–169, 2004. Preliminary version in STOC 2002.
- [47] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, Mar. 2002.
- [48] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.
- [49] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991. Preliminary version in CRYPTO 1989.
- [50] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the twenty-first annual ACM symposium on the theory of computing - STOC '89*, pages 33–43, Seattle, Washington, USA, May 1989. ACM.
- [51] P. Nguyen and J. Stern. Merkle-Hellman revisited: A cryptanalysis of the Qu-Vanstone cryptosystem based on group factorizations. In B. S. Kaliski, Jr., editor, *Advances in cryptology - CRYPTO '97, Proceedings of the 17th annual international cryptology conference*, volume 1294 of *Lecture Notes in Computer Science*, pages 198–212, Santa Barbara, California, USA, Aug. 1997. Springer.
- [52] P. Nguyen and J. Stern. Cryptanalysis of the Ajtai-Dwork cryptosystem. In H. Krawczyk, editor, *Advances in cryptology - CRYPTO '98, Proceedings of the 18th annual international cryptology conference*, volume 1462 of *Lecture Notes in Computer Science*, pages 223–242, Santa Barbara, California, USA, Aug. 1998. Springer.
- [53] P. Nguyen and J. Stern. Lattice reduction in cryptology: an update. In W. Bosma, editor, *Algorithmic number theory: 4th international symposium - ANTS-IV*, volume 1838 of *Lecture Notes in Computer Science*, pages 85–112, Leiden, The Netherlands, July 2000. Springer.
- [54] P. Nguyen and J. Stern. The two faces of lattices in cryptology. In J. Silverman, editor, *Cryptography and lattices conference - CaLC 2001*, volume 2146 of *Lecture Notes in Computer Science*, pages 146–180, Providence, RI, USA, Mar. 2001. Springer.
- [55] A. M. Odlyzko. The rise and fall of knapsack cryptosystems. In C. Pomerance, editor, *Cryptology and computational number theory*, volume 42 of *Proceedings of Symposia in Applied Mathematics*, pages 75–88, Boulder, Colorado, 1989. AMS.
- [56] G. Orton. A multiple-iterated trapdoor for dense compact knapsacks. In A. De Santis, editor, *Advances in cryptology - EUROCRYPT '94, Proceedings of a workshop on the theory and application of cryptographic techniques*, volume 950 of *Lecture Notes in Computer Science*, pages 112–130, Perugia, Italy, May 1994. Springer.
- [57] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In S. Halevi and T. Rabin, editors, *Theory of cryptography conference - Proceedings of TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166, New York, NY, USA, Mar. 2006. Springer.

- [58] O. Regev. Improved inapproximability of lattice and coding problems with preprocessing. *IEEE Transactions on Information Theory*, 50(9):2031–2037, 2004. Preliminary version in CCC 2003.
- [59] O. Regev. New lattice based cryptographic constructions. *Journal of the ACM*, 51(6):899–942, 2004. Preliminary version in STOC 2003.
- [60] H. Ritter. Breaking knapsack cryptosystems by max-norm enumeration. In *1st International Conference of the Theory and Applications of Cryptology - Pragocrypt '96*, pages 480–492, 1996.
- [61] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the twenty-second annual ACM symposium on the theory of computing - STOC '90*, pages 387–394, Baltimore, Maryland, USA, May 1990. ACM.
- [62] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53(2–3):201–224, 1987.
- [63] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66(1-3):181–199, Aug. 1994. Preliminary version in FCT 1991.
- [64] C.-P. Schnorr and H. H. Hörner. Attacking the Chor-Rivest cryptosystem by improved lattice reduction. In L. C. Guillou and J.-J. Quisquater, editors, *Advances in cryptology - EUROCRYPT '95, Proceedings of the international conference on the theory and application of cryptographic techniques*, volume 921 of *Lecture Notes in Computer Science*, pages 1–12, Saint-Malo, France, May 1995. Springer.
- [65] A. Shamir. A polynomial time algorithm for breaking the basic Merkle-Hellman cryptosystem. *IEEE Transactions on Information Theory*, 30(5):699–704, Sept. 1984. Preliminary version in FOCS'82.
- [66] M. Szydło. Hypercubic lattice reduction and analysis of GGH and NTRU signatures. In E. Biham, editor, *Advances in cryptology - EUROCRYPT 2003, proceedings of the international conference on the theory and application of cryptographic techniques*, volume 2656 of *Lecture Notes in Computer Science*, pages 433–448, Warsaw, Poland, May 2003. Springer.
- [67] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, Mathematische Instituut, University of Amsterdam, 1981. Available on-line at URL <http://turing.wins.uva.nl/~peter/>.